

A GAP TRICHOTOMY FOR BOOLEAN CONSTRAINT PROBLEMS: EXTENDING SCHAEFER’S THEOREM*

LUCY HAM

ABSTRACT. A computational problem exhibits a “gap property” when there is no tractable boundary between two disjoint sets of instances. We establish a Gap Trichotomy Theorem for a family of constraint problem variants, completely classifying the complexity of possible **NP**-hard gaps in the case of Boolean domains. As a consequence, we obtain a number of dichotomies for the complexity of specific variants of the constraint satisfaction problem: all are either polynomial-time tractable or **NP**-complete. Schaefer’s original dichotomy for SAT variants is a notable particular case.

Universal algebraic methods have been central to recent efforts in classifying the complexity of constraint satisfaction problems. A second contribution of the article is to develop aspects of the algebraic approach in the context of a number of variants of the constraint satisfaction problem. In particular, this allows us to lift our results on Boolean domains to many templates on non-Boolean domains.

1. INTRODUCTION

Constraint Satisfaction Problems (CSPs) occur widely in practice, both as natural problems, and as an underlying framework for constraint programming; see a text such as Tsang [37]. When the template is restricted to some fixed finite domain, these problems still cover many important practical problems as well as providing an important framework for theoretical considerations in computational complexity. In the case of Boolean (2-element) domains, constraint problems coincide with the SAT variants examined by Schaefer [33]. In his paper, Schaefer proved a famous dichotomy: he showed that the complexity of CSPs over a fixed Boolean constraint language is either decidable in polynomial time or is **NP**-complete. Since Schaefer’s seminal contribution, there have been enormous advances toward a more general dichotomy for constraint satisfaction problems on non-Boolean domains. In [16], Feder and Vardi argue that fixed template CSPs emerge as the broadest natural class for which a dichotomy might hold and propose the well-known *Dichotomy Conjecture*. Numerous extensions of Schaefer’s result are now known. The broadest of these include the case of three-element domains (Bulatov [9]), List Homomorphism Problems (Bulatov [10]), and the case of directed graphs without sources and sinks (Barto, Kozik and Niven [4]).

In addition to direct extensions of Schaefer’s results, many variants of constraint satisfaction problems have been shown to experience similar dichotomies, such as counting CSPs (Bulatov [11]), balanced CSPs (Schnoor and Schnoor [34])

Key words and phrases. Constraint Satisfaction Problem, Robust satisfiability, Clone theory, Dichotomy, Trichotomy, Boolean.

*This article is a heavily expanded version of an extended abstract submitted to the 2016 ISAAC proceedings, and differs by the inclusion of full proofs and most results of Section 8.

and equivalence problems (Böhler, Hemaspaandra, Reith and Vollmer [7]). In the present work, we explore computational complexity for variants of the CSP where instead of asking for the existence of a single solution, one asks for enough solutions to witness a range of conditions. We focus in particular on *separability* and *robust satisfiability*. The separation problem SEP asks if it is true that, for every pair of distinct variables u and v , there is a solution giving u a different value to v . The (k, \mathcal{F}) -robust satisfiability problem asks if every compatible partial assignment on k variables extends to a full solution. As explained in Jackson [21], the SEP condition arises naturally in universal algebraic considerations, and is also closely related to problems without a backbone: problems (typically SAT variants) where no variable is forced to take some fixed value. Implicit constraints such as these are widely associated with computational difficulty; see Monasson et al. [30], Beacham and Culberson [6] and Culberson and Gent [14]. In the language of [6] for example, the SEP condition corresponds to the “unfrozenness of equality”. Robust constraint problems have been studied in a number of different contexts, including Beacham [5] and Gottlob [19], with the general concept formulated in Abramsky, Gottlob and Kolaitis [1]. The (k, \mathcal{F}) -robustness condition is an extension of the robustness condition of [1].

A remarkable result of Gottlob [19] states that for any k , the decision problem $(3k + 3)\text{SAT}$ experiences a striking “gap property”: subject to $\mathbf{P} \neq \mathbf{NP}$, there is no polynomial-time decidable class of instances containing those for which every possible partial assignment on k variables extends to a satisfying solution and disjoint from the set of NO instances of the corresponding CSP. In fact, every such class is \mathbf{NP} -hard. Abramsky, Gottlob and Kolaitis [1] and then Jackson [21] showed that gaps of this style are also to be found for some other well-known \mathbf{NP} -complete problems, including 3SAT, G3C, NAE3SAT, and positive 1-in-3 SAT. We investigate computational gap theorems in the Boolean case, establishing a Gap Trichotomy Theorem (Theorem 4.2) that provides dichotomies for an entire family of computational problems, including SEP and $(2, \mathcal{F})$ -robust satisfiability, with all intractable cases exhibiting a gap property. In particular, we completely characterise the tractability of SEP and $(2, \mathcal{F})$ -robust satisfiability in the case of Boolean domains, as well recovering Schaefer’s Dichotomy Theorem in case of core relational structures.

A pivotal development in the classification of fixed template CSP complexity, including the aforementioned extensions of Schaefer’s dichotomy for Boolean CSPs, was the introduction of universal algebraic methods; starting with the work of Jeavons [24], Jeavons, Cohen, Gyssens [27], with the full framework presented in Bulatov, Jeavons, Krokhin [8]. The algebraic method concerns the analysis of “polymorphisms” (see Definition 3.2 below) of the template. For some computational problems, polymorphism analysis appears too coarse; this is true for SEP and robust satisfiability, as well as other variants of the CSP considered in [7] and [34]. In these cases, it is necessary to move to partial polymorphisms. As a second main result, we show that the basic universal algebraic methods can nevertheless be established in this setting, including many of the CSP variants just listed.

1.1. Organisation of article. The present article is comprised of eight sections: Section 2 introduces four computational problems that will be of primary interest. Section 3 covers the minimum required background for the full statement of our first main result. In Section 4, we give the necessary notion of a gap property and present

the Gap Trichotomy Theorem, as well as numerous dichotomy theorems that are immediate consequences. Section 5 covers preliminary concepts required for the main arguments; these relate to weak co-clones and strong partial clones. Sections 6 and 7 are dedicated to establishing the Gap Trichotomy Theorem; more specifically, Section 6 deals with the intractable cases, and Section 7 with the tractability cases. In Section 8, we investigate how algebraic methods for studying CSP complexity can be recovered for computational problems amenable to partial polymorphism analysis.

2. SEPARATION AND ROBUST SATISFIABILITY

We begin by introducing four computational problems that will be of primary focus in the present article.

Definition 2.1. Let Γ be a set of relation symbols, each with an associated finite arity. A *template* is a pair $\mathbb{A} = \langle A; \Gamma^{\mathbb{A}} \rangle$ consisting of a *finite* set A together with an interpretation of each n -ary relation symbol $r \in \Gamma$ as a subset $r^{\mathbb{A}}$ of A^n . The set $\Gamma^{\mathbb{A}} = \{r^{\mathbb{A}} \mid r \in \Gamma\}$ is often referred to as a *constraint language over domain A* .

When the context is clear, we will blur the distinction between a constraint language and the corresponding set of relation symbols by omitting the superscript \mathbb{A} .

Definition 2.2. Let $\mathbb{A} = \langle A; \Gamma^{\mathbb{A}} \rangle$ be a template. We define a *constraint instance* for $\Gamma^{\mathbb{A}}$ or a $\Gamma^{\mathbb{A}}$ -*instance* to be a triple $I = (V; A; \mathcal{C})$ consisting of

- a set of variables V ,
- the domain set A , and
- a set of constraints \mathcal{C} .

Each constraint $c \in \mathcal{C}$ is a pair $\langle s, r^{\mathbb{A}} \rangle$, where $r^{\mathbb{A}}$ is a k -ary relation in $\Gamma^{\mathbb{A}}$ and $s = (v_1, \dots, v_k)$ is a k -tuple involving variables from V . We define a *solution* of I to be any assignment $\phi: V \rightarrow A$ satisfying $(\phi(v_1), \dots, \phi(v_k)) \in r^{\mathbb{A}}$, for each $c = \langle (v_1, \dots, v_k), r^{\mathbb{A}} \rangle$ in \mathcal{C} .

Constraint satisfaction problem $\text{CSP}(\mathbb{A})$ over template \mathbb{A} .

Instance: a $\Gamma^{\mathbb{A}}$ -instance I .

Question: is there a solution of I ?

Nontrivial satisfaction problem $\text{CSP}_{\text{NonTriv}}(\mathbb{A})$ over template \mathbb{A} .

Instance: a $\Gamma^{\mathbb{A}}$ -instance I .

Question: is there a nonconstant solution of I ?

Separation problem $\text{SEP}(\mathbb{A})$ over template \mathbb{A} .

Instance: a $\Gamma^{\mathbb{A}}$ -instance I .

Question: for every pair $\{v_1, v_2\}$ of distinct variables in V , is there is a solution $\phi: V \rightarrow A$ of I such that $\phi(v_1) \neq \phi(v_2)$?

Note that every YES instance of a SEP problem is a YES instance of the corresponding CSP problem.

From an algebraic perspective, a constraint satisfaction problem over template $\mathbb{A} = \langle A; \Gamma^{\mathbb{A}} \rangle$ coincides with the homomorphism problem for the relational structure \mathbb{A} : an instance $I = (V; A; \mathcal{C})$ becomes a Γ -structure $\mathbb{I} = \langle V; \Gamma^{\mathbb{I}} \rangle$, by letting $r^{\mathbb{I}} = \{(v_1, \dots, v_k) \mid \langle (v_1, \dots, v_k), r^{\mathbb{A}} \rangle \in \mathcal{C}\}$, for each k -ary relation $r \in \Gamma$. If we let $\text{CSP}(\mathbb{A})$ be the class of all finite Γ -structures admitting a homomorphism into \mathbb{A} ,

then the constraint satisfaction problem over \mathbb{A} is equivalent to deciding membership in this class. In a similar way, we can view a separation problem $\text{SEP}(\mathbb{A})$ as deciding membership in the class $\text{RP}(\mathbb{A})$, which consists of all finite Γ -structures for which there is an injective homomorphism into a finite direct power of \mathbb{A} . See Jackson [21, §1] for a more detailed discussion.

Before we can introduce the fourth computational problem of interest, we require some definitions.

Definition 2.3. Let R be a set of finitary relation symbols and let $X = \{x_i \mid i \in I\}$ be a set of pairwise distinct variables. A formula in the language of R is called an *atomic formula* if it takes one of the following two forms:

- $x_i \approx x_j$ for some $x_i, x_j \in V$;
- $(x_1, \dots, x_k) \in r$, for some k -ary relation symbol $r \in R$ and some variables $x_1, \dots, x_k \in V$.

We refer to formulas of the second form *R-constraints* or *R-tuples*. A formula in the language of R is called a *primitive-positive formula* (abbreviated to pp-formula) if, for some $l \in \mathbb{N}_0$ and $m, n \in \mathbb{N}$, it is of the form:

$$(\exists w_1, \dots, w_\ell) \bigwedge_{i=1}^m \alpha_i(x_1, \dots, x_k, w_1, \dots, w_\ell),$$

where $w_1, \dots, w_\ell, x_1, \dots, x_k$ are distinct variables, and $\alpha_i(x_1, \dots, x_k, w_1, \dots, w_\ell)$ is an atomic formula with variables amongst $\{x_1, \dots, x_k, w_1, \dots, w_\ell\}$, for each $i \in \{1, \dots, m\}$. A quantifier-free pp-formula (corresponding to $\ell = 0$) is called a *conjunct-atomic formula*, and called a *equality-free conjunct-atomic formula* if it is a conjunct-atomic formula involving only *R-constraints*.

Definition 2.4. Let $\mathbb{A} = \langle A; \Gamma^{\mathbb{A}} \rangle$ be a template. Let $k \in \mathbb{N}$, let x_1, \dots, x_k be pairwise distinct variables and let $\rho(x_1, \dots, x_k)$ be a pp-formula in the language Γ . The formula ρ defines a k -ary relation $r_\rho^{\mathbb{A}}$ on A :

$$(a_1, \dots, a_k) \in r_\rho^{\mathbb{A}} \iff \rho(a_1, \dots, a_k) \text{ is true in } \mathbb{A}.$$

This extends to systems $(V; A; \mathcal{C})$ of variables constrained by relations in $\Gamma^{\mathbb{A}}$: include constraint $\langle (v_1, \dots, v_k), r_\rho^{\mathbb{A}} \rangle$ if there exist w_1, \dots, w_ℓ in V such that constraint $\langle (v_1, \dots, v_k, w_1, \dots, w_\ell), \alpha_i^{\mathbb{A}} \rangle$ is in \mathcal{C} for each $i \in \{1, \dots, m\}$. We refer to $\langle (v_1, \dots, v_k), r_\rho^{\mathbb{A}} \rangle$ as the *constraint defined by ρ from \mathcal{C}* .

Let \mathcal{F} be a finite set of pp-formulae in the language of Γ . We define the following constraint language over A :

$$\Gamma_{\mathcal{F}}^{\mathbb{A}} := \{r_\rho^{\mathbb{A}} \mid \rho \in \mathcal{F}\}.$$

We let $\mathbb{A}_{\mathcal{F}} = \langle A; \Gamma_{\mathcal{F}}^{\mathbb{A}} \rangle$ and for any instance $(V; A; \mathcal{C})$ of $\Gamma^{\mathbb{A}}$ we let $\mathcal{C}_{\mathcal{F}}$ be the constraints defined by \mathcal{F} from \mathcal{C} . We sometimes refer to the constraints in $\mathcal{C}_{\mathcal{F}}$ as *\mathcal{F} -constraints*.

Definition 2.5. Let $\Gamma^{\mathbb{A}}$ be a constraint language over a finite set A , let \mathcal{F} be a finite set of pp-formulae in the language of Γ and let $(V; A; \mathcal{C})$ be a $\Gamma^{\mathbb{A}}$ -instance. For a subset $S \subseteq V$, we say that an assignment $f: S \rightarrow A$ is *\mathcal{F} -compatible* if it preserves all constraints in $\mathcal{C}_{\mathcal{F}}$ on tuples from S , equivalently, if f is a solution of $(S; A; \mathcal{C}_{\mathcal{F}})$ in $\text{CSP}(\mathbb{A}_{\mathcal{F}})$.

In other words, if for some $\rho(x_1, \dots, x_k) \in \mathcal{F}$ and some tuple $(s_1, \dots, s_k) \in S^k$ the formula $\rho(s_1, \dots, s_k)$ is true in $(V; A; \mathcal{C})$, then $\rho(f(s_1), \dots, f(s_k))$ must be true in \mathbb{A} . In the following, we let k be a nonnegative integer and \mathcal{F} be a finite set of pp-formulae in the language of Γ .

The (k, \mathcal{F}) -robust satisfiability problem (k, \mathcal{F}) -Robust(\mathbb{A}) over template \mathbb{A} .

Instance: a $\Gamma^{\mathbb{A}}$ -instance I .

Question: does every \mathcal{F} -compatible assignment on k variables extend to a solution of I ?

In the case where \mathcal{F} consists of the set of pp-formulae defining all projections of relations in $\Gamma^{\mathbb{A}}$, the notion of \mathcal{F} -compatibility has been called “local compatibility” and (k, \mathcal{F}) -robust satisfiability called “ k -robust satisfiability”, see Abramsky, Gottlob and Kolaitis [1, §2], or [21, §3] for example. When \mathcal{F} consists of these formulae, we will use the abbreviated notation k -Robust(\mathbb{A}) to refer to the (k, \mathcal{F}) -robust satisfiability problem over \mathbb{A} . In [21, Lemma 3.1], Jackson proposes that \mathcal{F} -compatibility is a natural local compatibility condition.

We are primarily concerned with computational problems over finite constraint languages, however it will be useful to have a notion of tractability for constraint problems over infinite languages (over finite domains). The following definition is in line with Bulatov, Jeavons and Krokhin [8, Definition 2.7], though we use “locally tractable” in preference to “tractable”.

Definition 2.6. Let A be a finite non-empty set and let Γ be a set of relations on A , and let Prob be any of the computational problems CSP, CSP_{NonTriv}, SEP, or (k, \mathcal{F}) -Robust.

- If Γ is an infinite set, then we say that $\text{Prob}(\langle A; \Gamma \rangle)$ is *locally tractable* if, for every finite subset R of Γ , the problem $\text{Prob}(\langle A; R \rangle)$ is tractable.
- The problem $\text{Prob}(\langle A; \Gamma \rangle)$ is **NP**-complete if there is a finite subset $S \subseteq \Gamma$ for which $\text{Prob}(\langle A; S \rangle)$ is **NP**-complete.

We will sometimes abbreviate $\text{Prob}(\langle A; \Gamma \rangle)$ to $\text{Prob}(\Gamma)$.

3. BACKGROUND: POST’S LATTICE

The precise boundaries given in the Gap Trichotomy Theorem are more easily expressed using the language clone of theory. In this section, we introduce background material required to state the main results in full detail. The reader familiar with clone theory may skip this section and proceed to the main results in Section 4.

Definition 3.1. Let A be a non-empty set. For $n \in \mathbb{N}$, we define an *operation* of arity n on A to be a map $f: A^n \rightarrow A$. We define

\mathcal{R}_A to be the set of all non-empty, non-nullary finitary relations on A , and

\mathcal{O}_A to be the set of all non-empty, non-nullary finitary operations on A .

Definition 3.2. Let A be a non-empty set, let $n, k \in \mathbb{N}$, let $f: A^n \rightarrow A$ be a n -ary operation and let r be a k -ary relation on the set A . We say that f *preserves* r or r is *invariant under* f or f is a *polymorphism* of r , if whenever $a_1 = (a_{11}, \dots, a_{1n}), a_2 = (a_{21}, \dots, a_{2n}), \dots, a_k = (a_{k1}, \dots, a_{kn})$ are tuples in A^n , then

$$(\forall i \in \{1, \dots, n\} (a_{1i}, a_{2i}, \dots, a_{ki}) \in r) \implies (f(a_1), f(a_2), \dots, f(a_k)) \in r.$$

If F is set of operations then we say that r is invariant under F if r is invariant under every operation in F .

Definition 3.3. Let A be a non-empty set and let $\mathcal{C} \subseteq \mathcal{O}_A$. Then \mathcal{C} is a *clone* on the set A if the following two conditions hold:

- (1) \mathcal{C} contains all projection operations: that is, for all $n \in \mathbb{N}$, the i th projection $\pi_i: A^n \rightarrow A$ given by $\pi_i(x_1, \dots, x_n) = x_i$ belongs to \mathcal{C} ;
- (2) \mathcal{C} is closed under compositions.

For a set F of total operations, we can define $[F]$ to be the smallest clone containing F and we refer to $[F]$ as the *clone generated by F* . The set F is sometimes called a *base* for the clone $[F]$.

Definition 3.4. Let A be a non-empty set. A subset R of \mathcal{R}_A is called a *co-clone* or *relational clone* if it is closed under the formation of pp-definable relations. We can define $\langle R \rangle$ to be the smallest co-clone containing R and we refer to $\langle R \rangle$ as the *co-clone generated by R* . The set R is sometimes called a *base* for $\langle R \rangle$.

Recall that for any non-empty set A , the set of all clones (co-clones) on A ordered by set inclusion forms an algebraic intersection structure.

The sets $\wp(\mathcal{O}_A)$ and $\wp(\mathcal{R}_A)$ are complete lattices, where $\wp(-)$ is the powerset operator. The following well-known result was first observed by Geiger [17].

Theorem 3.5 ([17]). *Let A be a non-empty set. The following pair of maps $\text{Inv}: \wp(\mathcal{O}_A) \rightarrow \wp(\mathcal{R}_A)$ and $\text{Pol}: \wp(\mathcal{R}_A) \rightarrow \wp(\mathcal{O}_A)$ form a Galois connection between $\wp(\mathcal{O}_A)$ and $\wp(\mathcal{R}_A)$.*

$$\begin{aligned} \text{Inv}(F) &:= \{r \in \mathcal{R}_A \mid r \text{ is invariant under each } f \in F\} \text{ and} \\ \text{Pol}(R) &:= \{f \in \mathcal{O}_A \mid f \text{ preserves each } r \in R\}, \end{aligned}$$

for each $F \subseteq \mathcal{O}_A$ and each $R \subseteq \mathcal{R}_A$.

We list some useful consequences of Theorem 3.5.

- $\text{Pol}(\text{Inv}(F)) = [F]$, for all $F \subseteq \mathcal{O}_A$.
- $\text{Inv}(\text{Pol}(\Gamma)) = \langle \Gamma \rangle$, for all $\Gamma \subseteq \mathcal{R}_A$.
- The lattice of all clones on A is dually isomorphic to the lattice of all co-clones on A .
- Every clone \mathcal{C} corresponds to a unique co-clone $\text{Inv}(\mathcal{C})$.

Note that large clones correspond to small co-clones.

A characterisation of all clones on the two-element set $\{0, 1\}$ was given by Post [31] and the lattice of these “Boolean clones” is usually called *Post’s Lattice*. In fact, there is a countable infinity of clones (co-clones) in the Boolean case. An up-set of the Boolean co-clone lattice is given in Figure 1; the table included gives definitions of the shaded vertices in terms of relations invariant under basic operations. The operations c_0 and c_1 are the constant unary functions to 0 and 1, respectively, and \neg is the usual negation operation on $\{0, 1\}$. The shaded co-clones play a key role in formulating the main results; see Section 4 for a full discussion. We refer the reader to Schnoor [35, Table 3.1] for a full list of bases and definitions of all Boolean co-clones.

4. MAIN RESULTS

The first main result presents a trichotomy of computational gap theorems for Boolean constraint languages. As consequences, we obtain dichotomy theorems for

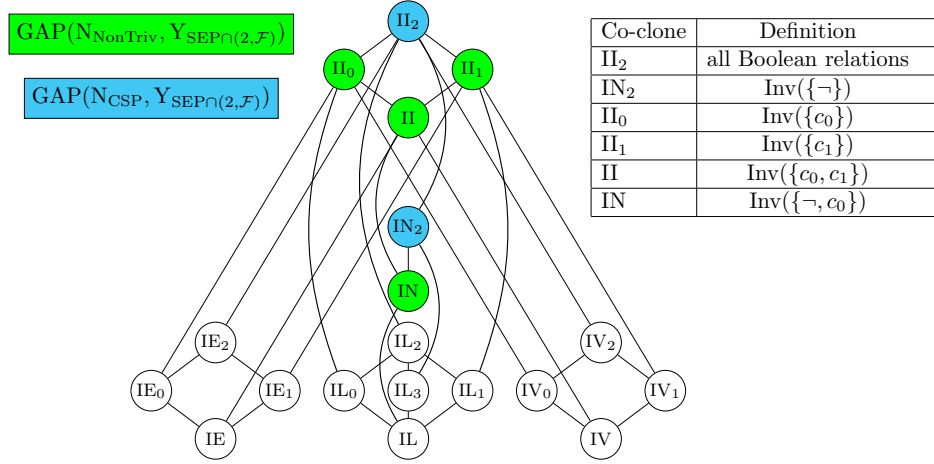


FIGURE 1. An upset in the Boolean co-clone lattice, with a table of polymorphism definitions for the shaded co-clones; IC abbreviates $\text{Inv}(\mathcal{C})$, for each Boolean clone \mathcal{C} .

each of the four computational problems described in Section 2. We first introduce the notion of a “gap property” and fix some useful notation.

In the following, let Γ be a finite set of relations on $\{0, 1\}$ and let P and Q be disjoint sets of Γ -instances.

The gap property $\text{GAP}(Q, P)$ for Γ with respect to Q and P .

$\text{GAP}(Q, P)$: any set of Γ -instances containing P and disjoint from Q has **NP**-hard membership with respect to polynomial-time Karp reductions.

Notation 4.1. Let Γ be a finite set of relations on $\{0, 1\}$, let $k \in \mathbb{N}$ and let \mathcal{F} be a finite set of pp-formulae in the language of Γ . We use

- $\text{N}_{\text{CSP}}(\Gamma)$ to denote the set of NO instances for $\text{CSP}(\Gamma)$,
- $\text{N}_{\text{NonTriv}}(\Gamma)$ to denote the set of NO instances for $\text{CSP}_{\text{NonTriv}}(\Gamma)$,
- $\text{Y}_{(k, \mathcal{F})}(\Gamma)$ to denote the set of YES instances for (k, \mathcal{F}) -Robust(Γ),
- $\text{Y}_{k\text{-Rob}}(\Gamma)$ to denote the set of YES instances for k -Robust(Γ),
- $\text{Y}_{\text{SEP}}(\Gamma)$ to denote the set of YES instances for $\text{SEP}(\Gamma)$,
- $\text{Y}_{\text{SEP} \cap (k, \mathcal{F})}(\Gamma)$ to denote the set of instances in $\text{Y}_{\text{SEP}}(\Gamma) \cap \text{Y}_{(2, \mathcal{F})}(\Gamma)$.

When the context refers to a specific constraint language Γ , we omit Γ from this notation.

Theorem 4.2 (Gap Trichotomy Theorem). *Let Γ be a constraint language on $\{0, 1\}$.*

- (1) *If $\langle \Gamma \rangle = \Pi_2$ or $\langle \Gamma \rangle = \text{IN}_2$, then Γ satisfies $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP}} \cap \text{Y}_{(2, \mathcal{F})})$, for some finite set \mathcal{F} of pp-formulae in Γ .*
- (2) *If $\langle \Gamma \rangle \in \{\Pi_1, \Pi_0, \Pi, \text{IN}\}$, then Γ satisfies $\text{GAP}(\text{N}_{\text{NonTriv}}, \text{Y}_{\text{SEP}} \cap \text{Y}_{(2, \mathcal{F})})$, for some finite set \mathcal{F} of pp-formulae in Γ .*
- (3) *In all other cases, the satisfiability problem $(2, \mathcal{F})$ -Robust(Γ) and the separation problem $\text{SEP}(\Gamma)$ are solvable in polynomial-time, for any finite set \mathcal{F} of pp-formulae in Γ .*

Shaded vertices in Figure 1 give the precise information for the Gap Trichotomy Theorem: constraint languages falling under the first item are coloured blue/dark grey in Figure 1. Those falling under the second item are green/light grey.

We now list four *dichotomy* theorems that are immediate consequences of the Gap Trichotomy Theorem. The Dichotomy Theorem for $\text{CSP}_{\text{NonTriv}}$ is a known result of Creignou and Hébrard [12, Proposition 3.5, 4.7].

Theorem 4.3 (Dichotomy Theorem for $(2, \mathcal{F})\text{-Robust}(\Gamma)$). *Let Γ be a constraint language on $\{0, 1\}$.*

- *If $\text{IN} \subseteq \langle \Gamma \rangle$ then $(2, \mathcal{F})\text{-Robust}(\Gamma)$ is **NP**-complete, for some finite set \mathcal{F} of pp-formulae in Γ .*
- *Otherwise, $(2, \mathcal{F})\text{-Robust}(\Gamma)$ is solvable in polynomial-time, for all finite sets \mathcal{F} of pp-formulae in Γ .*

Theorem 4.4 (Dichotomy Theorem for $\text{SEP}(\mathbb{A})$). *Let Γ be a constraint language on $\{0, 1\}$.*

- *If $\text{IN} \subseteq \langle \Gamma \rangle$ then $\text{SEP}(\Gamma)$ is **NP**-complete.*
- *Otherwise, $\text{SEP}(\Gamma)$ is solvable in polynomial-time.*

Theorem 4.5. [12, Dichotomy Theorem for $\text{CSP}_{\text{NonTriv}}(\mathbb{A})$] *Let Γ be a constraint language on $\{0, 1\}$.*

- *If $\text{IN} \subseteq \langle \Gamma \rangle$ then $\text{CSP}_{\text{NonTriv}}(\Gamma)$ is **NP**-complete.*
- *Otherwise, $\text{CSP}_{\text{NonTriv}}(\Gamma)$ is solvable in polynomial-time.*

In the following it is useful to recall that a finite relational structure (template) is a core if all of its unary polymorphisms are bijective.

4.1. Recovering Schaefer's Theorem. We observe that the intractability cases of Schaefer's Dichotomy Theorem are recovered from Theorem 4.2 in the case of core templates (Boolean CSPs over non-core templates are trivial). Indeed, when $\langle \{0, 1\}; \Gamma \rangle$ is a core template and $\text{IN} \subseteq \langle \Gamma \rangle$, we have $\langle \Gamma \rangle = \text{IN}_2$ or $\langle \Gamma \rangle = \text{II}_2$. By Theorem 4.2, it follows that Γ has gap property $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP} \cap (2, \mathcal{F})})$, for some finite set \mathcal{F} of pp-formulae in the language of Γ . As $\text{Y}_{\text{SEP}(2, \mathcal{F})} \subseteq \text{Y}_{\text{CSP}}$ and is disjoint from N_{CSP} , it follows that $\text{CSP}(\Gamma)$ is **NP**-complete.

Theorem 4.6 (Schaefer's Dichotomy Theorem [33]). *Let Γ be a constraint language on $\{0, 1\}$ such that $\langle \{0, 1\}; \Gamma \rangle$ is a core template.*

- *If $\text{IN} \subseteq \langle \Gamma \rangle$ then $\text{CSP}(\Gamma)$ is **NP**-complete.*
- *Otherwise, $\text{CSP}(\Gamma)$ is solvable in polynomial-time.*

5. WEAK CO-CLONES AND STRONG PARTIAL CLONES

We now give more technical definitions that are required for the main arguments.

Definition 5.1. For $n \in \mathbb{N}$, we define a *partial operation* of arity n on A to be a map $f: \text{dom}(f) \rightarrow A$, where $\text{dom}(f) \subseteq A^n$ and define \mathcal{P}_A to be the set of all non-empty, non-nullary finitary partial operations on A .

Definition 5.2. Let $n, k \in \mathbb{N}$, let $f: \text{dom}(f) \rightarrow A$ be a n -ary partial operation and let r be a k -ary relation on the set A . We say that f *preserves* r or r is *invariant under* f or f is a *partial polymorphism* of r , if whenever $a_1 = (a_{11}, \dots, a_{1n})$, $a_2 = (a_{21}, \dots, a_{2n})$, \dots , $a_k = (a_{k1}, \dots, a_{kn})$ are tuples in $\text{dom}(f)$, then the implication of Definition 3.2 holds. If F is set of partial operations then we say that r is invariant under F if r is invariant under every operation in F .

Definition 5.3. Let $m, n \in \mathbb{N}$, let $f \in \mathcal{P}_A$ be m -ary and let $g_1, \dots, g_m \in \mathcal{P}_A$ be n -ary. The *composition* $f(g_1, \dots, g_m)$ is an n -ary partial operation defined by

$$f(g_1, \dots, g_m)(x_1, \dots, x_n) := f(g_1(x_1, \dots, x_n), \dots, g_m(x_1, \dots, x_n)),$$

where $\text{dom}(f(g_1, \dots, g_m))$ is the set

$$\left\{ (x_1, \dots, x_n) \in \bigcap_{i=1}^m \text{dom}(g_i) \mid (g_1(x_1, \dots, x_n), \dots, g_m(x_1, \dots, x_n)) \in \text{dom}(f) \right\}.$$

Definition 5.4. Let $f, g \in \mathcal{P}_A$. We say that f is a *restriction* of g or that f *extends* to g if $\text{dom}(f) \subseteq \text{dom}(g)$ and f agrees with g on $\text{dom}(f)$.

Definition 5.5. Let A be a non-empty set and let $\mathcal{C} \subseteq \mathcal{P}_A$. Then \mathcal{C} is a *strong partial clone* if the following three conditions hold:

- (1) \mathcal{C} contains all total projection operations;
- (2) \mathcal{C} is closed under composition with non-empty domain;
- (3) \mathcal{C} is closed under arbitrary restriction of partial operations.

For a set F of partial operations, we can define $[F]_p$ to be the smallest strong partial clone containing F and we call $[F]_p$ the *strong partial clone generated* by F . The set F is sometimes called a *base* for the strong partial clone $[F]_p$.

Definition 5.6. Let A be a non-empty set. A subset R of \mathcal{R}_A is called a *weak co-clone* or *weak system* if it is closed under the formation of conjunct-atomic definable relations. We can define $\langle R \rangle_{\overline{\mathcal{A}}}$ to be the smallest weak co-clone containing R and we refer to $\langle R \rangle_{\overline{\mathcal{A}}}$ as the *weak co-clone generated* by R . The set R is sometimes called a *base* for the weak system $\langle R \rangle_{\overline{\mathcal{A}}}$.

A subset R of \mathcal{R}_A is called an *equality-free weak co-clone* or an *equality-free weak system* if it is closed under the formation of equality-free conjunct-atomic definable relations. We define $\langle R \rangle_{\overline{\mathcal{A}}, \neq}$ to be the smallest equality-free weak co-clone containing R and we refer to $\langle R \rangle_{\overline{\mathcal{A}}, \neq}$ as the *equality-free weak co-clone generated* by R . The set R is sometimes called a *base* for the equality-free weak system $\langle R \rangle_{\overline{\mathcal{A}}, \neq}$.

If we weaken the operators Inv and Pol , introduced in Section 3, to allow partial operations to be included in the definition, we obtain a refined Galois connection between the complete lattices $\wp(\mathcal{P}_A)$ and $\wp(\mathcal{R}_A)$. The Galois connection is finer in the sense that sets of relations are further separated on the basis of their expressibility power. The next theorem was first observed by Romov [32]. It can also be obtained as a special case of a Galois connection given in Davey, Pitkethly, and Willard [15, Section 1.2].

Theorem 5.7 ([17], [32]). *Let A be a non-empty set. The following pair of maps $\text{Inv}: \wp(\mathcal{P}_A) \rightarrow \wp(\mathcal{R}_A)$ and $\text{Pol}: \wp(\mathcal{R}_A) \rightarrow \wp(\mathcal{P}_A)$ form a Galois connection between the complete lattices $\wp(\mathcal{P}_A)$ and $\wp(\mathcal{R}_A)$.*

$$\text{Inv}(F) := \{r \in \mathcal{R}_A \mid r \text{ is invariant under each } f \in F\} \text{ and}$$

$$\text{pPol}(R) := \{f \in \mathcal{P}_A \mid f \text{ preserves each } r \in R\},$$

for each $F \subseteq \mathcal{P}_A$ and $R \subseteq \mathcal{R}_A$.

In a similar way that the Galois connection between operations and relations on a non-empty set A is connected to clones and co-clones, the Galois connection between partial operations and relations is closely related to strong partial clones and weak co-clones. In particular, sets of the form $\text{Inv}(F)$ are precisely the weak

co-clones and sets of the form $\text{pPol}(R)$ are precisely the strong partial clones, for $F \subseteq \mathcal{P}_A$ and $R \subseteq \mathcal{R}_A$. There are analogous consequences to those listed in Section 3, given for clones and co-clones.

The lattice of strong partial clones on $\{0, 1\}$ is complicated; it is uncountable and a complete classification appears difficult, see Schölzel [36] for example. Nevertheless, Post's lattice provides a useful approximation to lattice of strong partial clones in the Boolean setting: for each Boolean clone \mathcal{C} , it is known that the set of all strong partial clones whose total operations agree with \mathcal{C} forms an interval, and there are known generators for the top and bottom elements in each of these intervals.

In [34], Schnoor and Schnoor give a characterisation of the largest strong partial clone corresponding to a clone \mathcal{C} , and use this characterisation to give a canonical construction for finding a base for the smallest element of the associated interval in the weak co-clone lattice, see [34, Theorem 4.11] or [35, Theorem 3.9]. Such bases are referred to as *weak bases*, and from a complexity perspective, are crucial for establishing hardness results. On the other hand, generators for the largest elements of these intervals, referred to as *plain bases*, are in general key to proving tractability results; however an alternative approach is taken in this article. We mention only for completeness that plain bases for all Boolean co-clones can be found in Creignou, Kolaitis and Zanuttini [13, Table 2].

Definition 5.8. Let A be a non-empty set, let \mathcal{C} be a clone on A and let Γ be a set of finitary relations on A . We call Γ a *weak base* for the co-clone $\text{Inv}(\mathcal{C})$ if $\langle \Gamma \rangle_{\overline{\mathcal{C}}} = \text{Inv}(\mathcal{I}_{\cup}(\mathcal{C}))$, where $\mathcal{I}_{\cup}(\mathcal{C})$ is the largest strong partial clone corresponding to \mathcal{C} .

We will often present relations in a matrix form. The representation is not unique, but it is succinct.

Definition 5.9. Let $r = \{a_1, \dots, a_m\}$ be a k -ary relation on a non-empty set A with $|r| = m$. We define the matrix representation of r to be the $m \times k$ matrix $M = (a_{ij})$ over A with rows equal to the constraints of r .

Definition 5.10. Define Cols_3 to be the 8-ary relation over $\{0, 1\}$:

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 1 \end{bmatrix}$$

Definition 5.11. Let \mathcal{C} be clone on a non-empty set A and let r be a relation on A . Define $\mathcal{C}(r)$ to be the smallest relation containing r that is invariant under every operation in \mathcal{C} :

$$\mathcal{C}(r) := \bigcap \{s \mid s \in \text{Inv}(\mathcal{C}) \text{ and } r \subseteq s\}.$$

Following [34], we refer to $\mathcal{C}(r)$ as the \mathcal{C} -closure of the relation r and we say that r is a \mathcal{C} -core of $\mathcal{C}(r)$.

Using the work of Schnoor and Schnoor [34] and Schnoor [35, Table 3.1], the following construction gives weak-bases for each of the Boolean co-clones shaded in Figure 1.

Proposition 5.12. ([34, Theorem 4.11], [35, Table 3.1]) *Let IC be any of the Boolean co-clones listed in the table within Figure 1. Then $\mathcal{C}(\text{Cols}_3)$ is a weak-base for IC .*

Example 5.13. To construct a weak base for the Boolean co-clone $\text{IN}_2 = \text{Inv}(\{\neg\})$, we simply close the relation Cols_3 under \neg . Thus,

$$N_2(\text{Cols}_3) = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 1 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 1 & 1 & 0 \end{bmatrix}$$

Schnoor and Schnoor [34] observe that often the more restricted closure operator $\langle - \rangle_{\neg, \neq}$ fits more naturally to problems than the operator $\langle - \rangle_{\neg}$. This fact appears to be true of the computational problems considered in the present article. The following irredundancy condition ensures conjunct-atomic definability without equality.

Definition 5.14. [34, Definition 5.1] Let r be a k -ary relation on a non-empty set A . We say that the relation r is *=-redundant* if there exist a pair of equal columns in the matrix representation of r , and we say that r is \top -*redundant* if there is a relation s of arity $k-1$ such that the columns of r are exactly the columns of $s \times A$. A relation r is *irredundant* if it is neither =-redundant nor \top -redundant. A set Γ of finitary relations on A is *irredundant* if every relation r belonging to Γ is irredundant.

Remark 5.15. All six of the relations in Proposition 5.12 are easily seen to be irredundant.

Weak bases that generate not only the smallest weak-system, but also the smallest equality-free weak-system of relations generating the same co-clone, will be critical in establishing the main proofs to come.

Theorem 5.16. [34, Corollary 5.6] *Let A be a non-empty set, let \mathcal{C} be a clone on A and let Γ be an irredundant weak base for the co-clone $\text{Inv}(\mathcal{C})$. If Γ' is set of relations on A such that $\langle \Gamma' \rangle = \text{Inv}(\mathcal{C})$, then $\langle \Gamma \rangle_{\neg, \neq} \subseteq \langle \Gamma' \rangle_{\neg, \neq}$.*

6. TOWARDS A TRICHOTOMY: NP-GAPS

We now establish gap properties for relations generating the Boolean co-clones Π_2 , IN_2 , Π_0 , Π_1 , Π or IN ; recall exact definitions given in Figure 1. Each co-clone must be considered separately, however the proofs follow the same structure: we first establish a gap property for the irredundant weak-base and then use the fact that gap properties are preserved by the $\langle - \rangle_{\neg, \neq}$ operator.

We begin with three results that are crucial for establishing gap properties. The first theorem is an abridged version of [21, Theorem 6.1], and reveals that positive 1-in-3 SAT has the gap property $\text{GAP}(\text{N}_{\text{CSP}}, \mathcal{K})$, where \mathcal{K} is a particular subset of $\text{Y}_{2\text{-Rob}}$; see Theorem 6.1 below for the precise definition of X . Recall that 2 denotes the positive 1-in-3 SAT template $\langle \{0, 1\}; +1\text{in3 SAT} \rangle$.

Theorem 6.1. [21] *Let \mathcal{K} be the set consisting of all positive 1-in-3 SAT instances I with the following properties:*

- *no variable appears more than once in each constraint tuple of I ,*
- *I is 2-robustly positive 1-in-3 satisfiable.*

Any set of positive 1-in-3 SAT instances containing \mathcal{K} and contained within the members of $\text{CSP}(2)$ has **NP**-hard membership with respect to polynomial-time Karp reductions.

The next lemma summarises the basic method employed in Abramsky, Gottlob and Kolaitis [1] and Jackson [21]. It is essentially the definition of reduction for promise problems; see [18, Definition 3], for example.

Lemma 6.2. *Let Γ and Γ' be finite sets of relations on $\{0, 1\}$. Let A and B be disjoint sets of Γ -instances and let X and Y be disjoint sets of Γ' -instances. Further, let Γ have the gap property $\text{GAP}(A, B)$. If there is a polynomial-time computable function $f: \mathcal{I}_\Gamma \rightarrow \mathcal{I}_{\Gamma'}$ satisfying:*

- (1) $I \in A \Rightarrow f(I) \in X$,
- (2) $I \in B \Rightarrow f(I) \in Y$,

then Γ' has the gap property $\text{GAP}(X, Y)$. In particular, Γ' has the gap property $\text{GAP}(f(A), f(B))$

It is well known that the complexity of $\text{CSP}(\Gamma)$ depends only on the co-clone generated by Γ , see [24, Theorem 3.4] or alternatively [8, Theorem 2.16] for a proof explicitly using pp-formulae. We give an analogous result relating equality-free conjunct atomic definitions to reductions for SEP and $(2, \mathcal{F})$ -robust satisfiability. This result is stated below as Theorem 6.4 and will be crucial for establishing hardness results in this section. We first introduce a lemma that will help streamline the proof.

Lemma 6.3. *Let $\Gamma^\mathbb{A}$ be a constraint language over a finite set A and let $R^\mathbb{A}$ be a finite set of relations in $\langle \Gamma^\mathbb{A} \rangle_{\exists, \neq}$. There is a polynomial-time construction that transforms any instance $I = (V; A; \mathcal{C})$ of $\text{CSP}(R^\mathbb{A})$ into an instance I' of $\text{CSP}(\Gamma^\mathbb{A})$, and moreover, the solutions of I are exactly the solutions of I' .*

Proof. The following construction is borrowed from [8, Theorem 2.16]. Given an instance $I = (V; A; \mathcal{C})$ of $\text{CSP}(\langle A; R^\mathbb{A} \rangle)$, construct an instance $I' = (V; A; \mathcal{C}')$ of $\text{CSP}(\langle A; \Gamma^\mathbb{A} \rangle)$ in the following way. For each constraint $\langle (v_1, \dots, v_\ell), r^\mathbb{A} \rangle \in \mathcal{C}$ with $r^\mathbb{A}$ representable by the equality-free conjunct-atomic formula:

$$r(v_1, \dots, v_\ell) = s_1(w_1^1, \dots, w_{\ell_1}^1) \wedge \dots \wedge s_n(w_1^n, \dots, w_{\ell_n}^n),$$

where $\{w_1^1, \dots, w_{\ell_1}^1, \dots, w_1^n, \dots, w_{\ell_n}^n\} \subseteq \{v_1, \dots, v_\ell\}$, and $\{s_1, \dots, s_n\} \subseteq \Gamma$, we add the constraints $c_1 = \langle (w_1^1, \dots, w_{\ell_1}^1), s_1^\mathbb{A} \rangle, \dots, c_n = \langle (w_1^n, \dots, w_{\ell_n}^n), s_n^\mathbb{A} \rangle$ to \mathcal{C} , and remove $\langle (v_1, \dots, v_\ell), r^\mathbb{A} \rangle$ from \mathcal{C} . Since $r^\mathbb{A}(a_1, \dots, a_\ell)$ is true in \mathbb{A} if and only if $(a_1^1, \dots, a_{\ell_1}^1) \in s_1^\mathbb{A}, \dots, (a_1^n, \dots, a_{\ell_n}^n) \in s_n^\mathbb{A}$, for $(a_1, \dots, a_\ell) \in A^\ell$ and $r^\mathbb{A} \in R^\mathbb{A}$, it follows that the solutions of I are precisely the solutions of I' . \square

Theorem 6.4. *Let $\Gamma^\mathbb{A}$ be a constraint language over a set A , let $R^\mathbb{A}$ be any finite set of relations in $\langle \Gamma^\mathbb{A} \rangle_{\exists, \neq}$, let \mathcal{F} be a finite set of pp-formulae in the language of R and let $k \in \mathbb{N}$. There is a polynomial-time computable function that reduces*

- (1) $\text{CSP}(R^\mathbb{A})$ to $\text{CSP}(\Gamma^\mathbb{A})$,
- (2) $\text{SEP}(R^\mathbb{A})$ to $\text{SEP}(\Gamma^\mathbb{A})$, and
- (3) $(k, \mathcal{F})\text{-Robust}(R^\mathbb{A})$ to $(k, \mathcal{G})\text{-Robust}(\Gamma^\mathbb{A})$, for some finite set \mathcal{G} of pp-formulae in the language of Γ .

Proof. We transform an R -instance $I = (V; A; \mathcal{C})$ into a Γ -instance $I' = (V; A; \mathcal{C}')$ according to the construction given in the proof of Lemma 6.3. The reduction from $\text{CSP}(R^{\mathbb{A}})$ to $\text{CSP}(\Gamma^{\mathbb{A}})$ is then obtained immediately from Lemma 6.3. This proves (1).

We now prove (2). Since the solutions of I in $\text{CSP}(R^{\mathbb{A}})$ are precisely the solutions of I' in $\text{CSP}(\Gamma^{\mathbb{A}})$, it follows that separating solutions of I are exactly the separating solutions of I' . Hence I is a YES instance of $\text{SEP}(R^{\mathbb{A}})$ if and only if I' is a YES instance of $\text{SEP}(\Gamma^{\mathbb{A}})$.

For (3), consider $r \in R$ of arity ℓ and abstractly expressible by an equality-free conjunct-atomic formula $r(x_1, \dots, x_\ell)$ in the language of Γ . For each pp-formula $\rho(w_1, \dots, w_m) \in \mathcal{F}$, we construct a pp-formula $\rho_\Gamma(w_1, \dots, w_m)$ in the language of Γ in the following way: replace every occurrence of an ℓ -ary relation symbol r in ρ with its conjunct-atomic defining formula $r(x_1, \dots, x_\ell)$. Let $\mathcal{G} = \{\rho_\Gamma \mid \rho \in \mathcal{F}\}$. Then since $\rho(a_1, \dots, a_m)$ is true in $\langle A; R^{\mathbb{A}} \rangle$ if and only if $\rho_\Gamma(a_1, \dots, a_m)$ is true in $\langle A; \Gamma^{\mathbb{A}} \rangle$ for $(a_1, \dots, a_m) \in A^m$ and $\rho(w_1, \dots, w_m) \in \mathcal{F}$, it follows that the \mathcal{F} -compatible assignments on k variables of I are exactly the \mathcal{G} -compatible assignments on k variables of I' . Thus, since the solutions of I are precisely the solutions of I' , by Lemma 6.3, it then follows that I is a YES instance of $(k, \mathcal{F})\text{-Robust}(R^{\mathbb{A}})$ if and only if I' is a YES instance of $(k, \mathcal{G})\text{-Robust}(\Gamma^{\mathbb{A}})$. \square

Theorem 6.4 holds more generally: with some caveats and proper amendment to the proof, the assumption that $R \subseteq \langle \Gamma \rangle_{\neq, \neq}$ can be weakened to $R \subseteq \langle \Gamma \rangle_{\neq}$. However, this result is not required for establishing our main theorems. We refer the reader to Section 8 for a proof.

For the purpose of easy reference, we extract the following consequence of Theorem 6.4.

Corollary 6.5. *Let Γ be a set of relations on $\{0, 1\}$, let R be a finite set of relations in $\langle \Gamma \rangle_{\neq, \neq}$ and let \mathcal{F} be a finite set of pp-formulae in the language of R . There is a polynomial-time computable function taking*

- $\text{N}_{\text{CSP}}(R)$ to $\text{N}_{\text{CSP}}(\Gamma)$, and
- $\text{Y}_{\text{SEP} \cap (2, \mathcal{F})}(R)$ to $\text{Y}_{\text{SEP} \cap (2, \mathcal{G})}(\Gamma)$, for some finite set \mathcal{G} of pp-formulae in Γ .

The first two subsections cover the co-clones II_2 and IN_2 corresponding to Statement 1 of the Gap Trichotomy Theorem.

6.1. The Boolean co-clone II_2 . Recall the relation Cols_3 given in Definition 5.10. By Proposition 5.12, the relation $\text{I}_2(\text{Cols}_3) = \text{Cols}_3$, is a weak base for II_2 . We will use III_2 to denote the template $\langle \{0, 1\}; \text{Cols}_3 \rangle$ and $\text{II}_2\text{-SAT}$ as notation for the constraint problem $\text{CSP}(\text{III}_2)$.

In the following, it is useful to remember that the positive 1-in-3 SAT relation has gap property $\text{GAP}(\text{N}_{\text{CSP}}, \mathcal{K})$, where \mathcal{K} satisfies the properties given in Theorem 6.1.

Proposition 6.6. *The relation $\text{I}_2(\text{Cols}_3)$ satisfies $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP} \cap 2\text{-Rob}})$.*

Proof. We begin by showing there is a polynomial-time Karp reduction from $\text{CSP}(2)$ to $\text{II}_2\text{-SAT}$. We then show that this reduction takes instances with the properties given in Theorem 6.1 to instances in $\text{Y}_{\text{SEP} \cap 2\text{-Rob}}(\text{I}_2(\text{Cols}_3))$. The result will then follow from Theorem 6.1 and Lemma 6.2.

We motivate the following construction by observing that columns 4, 5 and 6 of $I_2(\text{Cols}_3)$ are the negations of columns 1, 2 and 3, respectively; refer to Definition 5.10. Given an instance $I = (V; \{0, 1\}; \mathcal{C})$ of positive 1-in-3 SAT, we construct an instance $I^* = (V^*; \{0, 1\}; \mathcal{C}^*)$ of Π_2 -SAT in the following way.

- (i) First let $\bar{V} = \{\bar{v} \mid v \in V\}$ be a disjoint copy of V , and then construct $V^* = V \cup \bar{V} \cup \{\top, \perp\}$, where $\top, \perp \notin V \cup \bar{V}$,
- (ii) for each constraint $\langle (x, y, z), +1\text{in3SAT} \rangle$ in \mathcal{C} , we add the constraint $\langle (x, y, z, \bar{x}, \bar{y}, \bar{z}, \perp, \top), I_2(\text{Cols}_3) \rangle$ to \mathcal{C} and remove $\langle (x, y, z), +1\text{in3SAT} \rangle$.

Any solution φ of I in $\text{CSP}(2)$ can be extended to a solution φ^* of I^* in the following way. For each $v \in V$, define $\varphi^*(v) := \varphi(v)$, $\varphi^*(\bar{v}) = \neg \circ \varphi(v)$, $\varphi^*(\perp) = 0$ and $\varphi^*(\top) = 1$, where \neg is the usual Boolean complement. For the converse direction, observe that the projection $\pi_{\{1,2,3\}}(I_2(\text{Cols}_3)) = +1\text{in3SAT}$, thus if ψ is a solution of I^* in Π_2 -SAT, then the restriction $\psi|_V$ is a solution of I in $\text{CSP}(2)$. Hence we have shown that any solution ϕ of I extends uniquely to a solution ϕ^* of I^* .

Now assume that I has the properties listed in Theorem 6.1. Since I is 2-robustly satisfiable and no variable appears more than once in each constraint tuple, it follows that I has the following properties.

- (\heartsuit) For every pair of distinct variables x and y in V , there are solutions φ_1, φ_2 and φ_3 of I satisfying

$$\begin{aligned} (\varphi_1(x), \varphi_1(y)) &= (0, 0), \\ (\varphi_2(x), \varphi_2(y)) &= (0, 1), \\ (\varphi_3(x), \varphi_3(y)) &= (1, 0). \end{aligned}$$

- (\diamond) If x and y do not appear in a common constraint tuple, then there is a solution φ_4 of I satisfying $(\varphi_4(x), \varphi_4(y)) = (1, 1)$.

We first show that $I^* \in \mathbf{Y}_{2\text{-Rob}}(I_2(\text{Cols}_3))$. Let $\{u, v\}$ be a pair of distinct variables in V^* and let $\alpha: \{u, v\} \rightarrow \{0, 1\}$ be an assignment that is locally compatible for I^* . There are four cases to consider.

Case 1. Assume that $\{u, v\} \subseteq V$.

Subcase 1(a): If u and v appear in a common constraint tuple, then there are three locally compatible assignments to check: α satisfying

$$(\alpha(u), \alpha(v)) = (0, 0), (\alpha(u), \alpha(v)) = (0, 1), \text{ or } (\alpha(u), \alpha(v)) = (1, 0).$$

From (\heartsuit), there is a solution of I extending each of these assignments, namely φ_1, φ_2 and φ_3 , respectively. Using the fact that every solution of I extends to a solution of I^* , we get solutions φ_1^*, φ_2^* and φ_3^* of I^* which extend in the required ways.

Subcase 1(b): If u and v do not appear in a common constraint tuple, then every possible assignment for $\{u, v\}$ is locally compatible. Now apply (\heartsuit) and (\diamond) and the fact that every solution I extends to a solution of I^* , to get solutions of I^* for every possible assignment.

Case 2. Assume that $u \in V$ and $v \in \bar{V}$. We will consider cases only up to symmetry.

Subcase 2(a): If $v = \bar{u}$. In this case there two locally compatible assignments for $\{u, v\}$: $(\alpha(u), \alpha(v)) = (0, 1)$ or $(\alpha(u), \alpha(v)) = (1, 0)$. Consider first α satisfying $(\alpha(u), \alpha(v)) = (0, 1)$. Using (\heartsuit), there is a solution φ_2 of I extending this assignment. Since φ_2 satisfies $\varphi_2(u) = 0$, it follows that $\varphi_2^*(u) = 0$ and $\varphi_2^*(v) = \varphi_2^*(\bar{u}) = \neg \circ \varphi_2(u) = 1$. Hence φ_2^* is a solution of I^* extending this assignment. The case for $(\alpha(u), \alpha(v)) = (1, 0)$ is similar.

Subcase 2(b): If $v = \bar{w}$, for some $w \in V$ with $w \neq u$, then there are two subcases to check.

2(b)(i): If u and w (and hence v) appear in a common constraint tuple. There are three locally compatible assignments for $\{u, v\}$ to check in this case: α satisfying

$$(\alpha(u), \alpha(v)) = (0, 1), (\alpha(u), \alpha(v)) = (1, 1), \text{ or } (\alpha(u), \alpha(v)) = (0, 0).$$

Consider first $(\alpha(u), \alpha(v)) = (0, 1)$. From (\heartsuit) , there is a solution φ_1 of I satisfying $(\varphi_1(x), \varphi_1(y)) = (0, 0)$. It then follows that $\varphi_1^*(u) = 0$ and $\varphi_1^*(v) = \varphi_1^*(\bar{w}) = \neg \circ \varphi_1(w) = 1$. Hence we have found a solution of I^* extending this assignment. Now consider $(\alpha(u), \alpha(v)) = (1, 1)$. There is a solution φ_3 of I satisfying $(\varphi_3(u), \varphi_3(w)) = (1, 0)$, guaranteed by (\heartsuit) . It follows that $\varphi_3^*(u) = 1$ and $\varphi_3^*(v) = \varphi_3^*(\bar{w}) = 1$. In the case where α satisfies $(\alpha(u), \alpha(v)) = (0, 0)$, the solution φ_2^* gives the required extension.

2(b)(ii): If u and v do not appear in a common constraint tuple, then every possible assignment for $\{u, v\}$ is locally compatible. For the assignments $(\alpha(u), \alpha(v)) = (0, 1)$, $(\alpha(u), \alpha(v)) = (1, 1)$ and $(\alpha(u), \alpha(v)) = (0, 0)$, we apply Subcase 2(b)(i). It remains to show there is a solution extending $(\alpha(u), \alpha(v)) = (1, 0)$. From (\diamond) , there is a solution φ_4 of I satisfying $(\varphi_4(u), \varphi_4(w)) = (1, 1)$. It follows that $\varphi_4^*(u) = 1$ and $\varphi_4^*(w) = \varphi_4^*(\bar{w}) = \neg \circ \varphi_4(w) = 0$.

Case 3. Assume that $\{u, v\} \in \bar{V}$. Then $u = \bar{z}$ for some $z \in V$ and $v = \bar{w}$, for some $w \in V$. Apply Case 1 to the pair $\{z, w\}$.

Case 4: If $v \in \{\perp, \top\}$ then there are three subcases to consider.

Subcase 4(a): If $u \in \{\perp, \top\}$, then up to symmetry the only locally compatible assignment is $(\alpha(u), \alpha(v)) = (0, 1)$. A solution of I^* extending this assignment is obtained from any solution of I .

Subcase 4(b): If $u \in V$, then up to symmetry in the second coordinate there are two locally compatible assignments to consider: α satisfying $(\alpha(u), \alpha(v)) = (0, 0)$ or $(\alpha(u), \alpha(v)) = (1, 0)$. For the assignment $(\alpha(u), \alpha(v)) = (0, 0)$, use any solution φ of I satisfying $\varphi(u) = 1$ guaranteed by (\heartsuit) ; since the extension φ^* of any solution of I satisfies $\varphi^*(\perp) = 0$. For $(\alpha(u), \alpha(v)) = (1, 0)$ use any solution φ of I satisfying $\varphi(u) = 1$ guaranteed by (\heartsuit) .

Subcase 4(b): If $u = \bar{w}$, for some $w \in V$. In this case, we apply Subcase 4(b) to the pair $\{w, v\}$.

To show $I^* \in \text{Y}_{\text{SEP}}(\text{I}_2(\text{Cols}_3))$ observe that Cases 1 and 2 above show that for any pair of distinct variables $\{u, v\}$ in V^* , there is a solution of I^* in $\text{CSP}(\text{III}_2)$ separating u and v . \square

Remark 6.7. It is immediate from the proof that there is no loss of generality if the following extra condition is included in Proposition 6.6 as an allowed assumption on instances: there are variables \perp and \top , such that in every constraint tuple, the final two coordinates are \perp and \top , in that order.

The next result follows immediately from Proposition 6.6.

Corollary 6.8. *The following problems are NP-complete.*

- *The 2-robust satisfiability problem 2-Robust(III₂).*
- *The separation problem SEP(III₂).*

We are now ready to give the main result for this subsection.

Theorem 6.9. *Let Γ be a constraint language on $\{0, 1\}$ such that $\langle \Gamma \rangle = \Pi_2$. Then Γ has the gap property $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP} \cap (2, \mathcal{G})})$, for some finite set \mathcal{G} of pp-formulae in the language of Γ , and consequently both $(2, \mathcal{G})\text{-Robust}(\Gamma)$ and $\text{SEP}(\Gamma)$ are **NP**-complete.*

Proof. Since $\text{I}_2(\text{Cols}_3)$ is an irredundant weak-base for the co-clone Π_2 , by Theorem 5.16 we have $\text{I}_2(\text{Cols}_3) \in \langle \Gamma \rangle_{\neq}$. From Proposition 6.6, we have $\text{I}_2(\text{Cols}_3)$ satisfies $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP} \cap 2\text{-Rob}})$. Now applying Corollary 6.5, it follows from Lemma 6.2, that Γ has the gap property $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP} \cap (2, \mathcal{G})})$, for some finite set \mathcal{G} of pp-formulae in the language of Γ , and consequently $(2, \mathcal{G})\text{-Robust}(\Gamma)$ and $\text{SEP}(\Gamma)$ are **NP**-complete. \square

6.2. The Boolean co-clone IN_2 . Recall that the relation $\text{N}_2(\text{Cols}_3)$ constructed in Example 5.13 is a weak base for IN_2 . We will use \mathbb{IN}_2 to denote the corresponding template $\langle \{0, 1\}; \text{N}_2(\text{Cols}_3) \rangle$ and $\text{IN}_2\text{-SAT}$ as notation for the constraint problem $\text{CSP}(\mathbb{IN}_2)$.

Proposition 6.10. *The relation $\text{N}_2(\text{Cols}_3)$ satisfies $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP} \cap 2\text{-Rob}})$.*

Proof. We show there is polynomial-time Karp reduction from positive 1-in-3 SAT to $\text{IN}_2\text{-SAT}$ that takes instances satisfying the properties listed in Theorem 6.1 to instances in $\text{Y}_{\text{SEP} \cap 2\text{-Rob}}(\text{N}_2(\text{Cols}_3))$. The result then follows from Theorem 6.1 and Lemma 6.2.

Given an instance $I = (V; \{0, 1\}; \mathcal{C})$ of positive 1-in-3 SAT, we construct an instance $I^\# = (V^\#; \{0, 1\}; \mathcal{C}^\#)$ of $\text{IN}_2\text{-SAT}$ in the following way:

- (i) for each variable $v \in V$, we add the variable \bar{v} to V (each of the variables \bar{v} is a new variable not in V),
- (ii) add the new variables \top and \perp to V ,
- (iii) for each constraint $\langle (x, y, z), +1\text{in3 SAT} \rangle$ in \mathcal{C} , we add the constraints

$$\langle (x, y, z, \bar{x}, \bar{y}, \bar{z}, \perp, \top), \text{N}_2(\text{Cols}_3) \rangle, \text{ and } \\ \langle (\bar{x}, \bar{y}, \bar{z}, x, y, z, \top, \perp), \text{N}_2(\text{Cols}_3) \rangle$$

to \mathcal{C} and remove $\langle (x, y, z), +1\text{in3 SAT} \rangle$ from \mathcal{C} .

Any solution φ of I in $\text{CSP}(2)$ can be extended to a solution $\varphi^\#$ of $I^\#$ in $\text{IN}_2\text{-SAT}$ in the same way as in Proposition 6.6: for all $v \in V$, define $\varphi^\#(v) := \varphi(v)$, $\varphi^\#(\bar{v}) := \neg \circ \varphi(v)$, $\varphi^\#(\perp) := 0$ and $\varphi^\#(\top) := 1$. For the converse direction, we will show that for any solution $\varphi^\#$ of $I^\#$, either $\varphi^\#$ or $\neg \circ \varphi^\#$ can be restricted to a solution of I . Since the map $\neg : \{0, 1\} \rightarrow \{0, 1\}$ is an automorphism of \mathbb{IN}_2 , it follows that $\neg \circ \varphi^\#$ is solution of $I^\#$, whenever $\varphi^\#$ is a solution of $I^\#$. Thus, without loss of generality we may assume that $\varphi^\#(\perp) = 0$. This implies that $\varphi^\#$ satisfies constraints of the form $\langle (x, y, z, \bar{x}, \bar{y}, \bar{z}, \perp, \top), \text{N}_2(\text{Cols}_3) \rangle$ into the subset $\text{I}_2(\text{Cols}_3)$ of $\text{N}_2(\text{Cols}_3)$ (that is, the first three rows of $\text{N}_2(\text{Cols}_3)$). Then $\varphi^\#|_V$ is a solution of I for positive 1-in-3 SAT.

Now assume that I has the properties given in Theorem 6.1. We show that $I^\# \in \text{Y}_{2\text{-Rob}}(\text{N}_2(\text{Cols}_3))$ by showing that every locally compatible assignment on two variables extends to a solution of $I^\#$. Recall the instance $I^* = (V^*; \{0, 1\}; \mathcal{C}^*)$ constructed in Proposition 6.6. The following fact will be useful.

- (\star) for every solution φ^* of I^* in $\text{IN}_2\text{-SAT}$, both φ^* and $\neg \circ \varphi^*$ is a solution of $I^\#$ in $\text{IN}_2\text{-SAT}$.

Let $\{u, v\}$ be distinct variables in V and let $\alpha: \{u, v\} \rightarrow \{0, 1\}$ be an assignment that is locally compatible for $I^\#$. There are two cases to consider.

Case 1. If α is locally compatible assignment for I^\star , then use the fact that I^\star is in $Y_{2\text{-Rob}}(I_2(\text{Cols}_3))$ to extend α to a solution ψ^\star of I^\star , which is also a solution of $I^\#$ by (\star) .

Case 2. If α is not locally compatible for I^\star then it is straightforward to check that $\neg \circ \alpha$ is a locally compatible for I^\star . We then use the fact that I^\star is in $Y_{2\text{-Rob}}(I_2(\text{Cols}_3))$ to extend $\neg \circ \alpha$ to a solution ϕ^\star of I^\star . Then $\neg \circ \phi^\star$ is a solution of $I^\#$ by (\star) and extends α .

It remains to show $I^\# \in Y_{\text{SEP}}(N_2(\text{Cols}_3))$. Let $\{u, v\}$ be any pair of distinct variables in $V^\#$. Then $\{u, v\} \subseteq V^\star$. Since $I^\star \in Y_{\text{SEP}}(I_2(\text{Cols}_3))$, there is a solution φ^\star of I^\star separating u and v . Then φ^\star is a solution of $I^\#$ that separates u and v , by (\star) . \square

Remark 6.11. As with Remark 6.7, we may additionally add the following technical assumption to Proposition 6.10: there are two variables \perp and \top and every constraint tuple has its final two coordinates equal to either \perp, \top or \top, \perp .

The next result immediately follows from Proposition 6.10.

Corollary 6.12. *The following problems are NP-complete.*

- *The 2-robust satisfiability problem $2\text{-Robust}(\mathbb{IN}_2)$.*
- *The separation problem $\text{SEP}(\mathbb{IN}_2)$.*

We state the main result for this subsection.

Theorem 6.13. *Let Γ be a constraint language on $\{0, 1\}$ such that $\langle \Gamma \rangle = \mathbb{IN}_2$. Then Γ has the gap property $\text{GAP}(N_{\text{CSP}}, Y_{\text{SEP} \cap (2, \mathcal{G})})$, for some finite set \mathcal{G} of pp-formulae in the language of Γ , and consequently the problems $(2, \mathcal{G})\text{-Robust}(\Gamma)$ and $\text{SEP}(\Gamma)$ are NP-complete.*

Proof. It is easily verified that the weak base $N_2(\text{Cols}_3)$ is irredundant. By Theorem 5.16, it follows that $N_2(\text{Cols}_3) \in \langle \Gamma \rangle_{\neq}$. From Proposition 6.10, we have $N_2(\text{Cols}_3)$ satisfies $\text{GAP}(N_{\text{CSP}}, Y_{\text{SEP} \cap 2\text{-Rob}})$. Applying Corollary 6.5, it follows from Lemma 6.2, that Γ has the gap property $\text{GAP}(N_{\text{CSP}}, Y_{\text{SEP} \cap (2, \mathcal{G})})$, for some finite set \mathcal{G} of pp-formulae in the language of Γ , and consequently $(2, \mathcal{G})\text{-Robust}(\Gamma)$ and $\text{SEP}(\Gamma)$ are NP-complete. \square

For the remaining clones $\mathbb{II}_1, \mathbb{II}_0, \mathbb{II}, \mathbb{IN}$ corresponding to Statement 2 in the Gap Dichotomy Theorem 4.2, we can reuse the same fundamental construction used for \mathbb{II}_2 and \mathbb{IN}_2 . In all cases, the proof proceeds as follows: to lie in Y_{SEP} or $Y_{(2, \mathcal{F})}$, it is necessary to have an assignment in which \perp and \top take different values. We then argue that this forces solutions into \mathbb{II}_2 or \mathbb{IN}_2 .

6.3. The Boolean co-clone \mathbb{II}_1 . From Proposition 5.12, the relation $I_1(\text{Cols}_3)$ is a weak base for the co-clone \mathbb{II}_1 :

$$I_1(\text{Cols}_3) = \begin{bmatrix} 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}$$

We will use \mathbb{III}_1 to denote the template $\langle \{0, 1\}; I_1(\text{Cols}_3) \rangle$ and $\mathbb{II}_1\text{-SAT}$ to denote $\text{CSP}(\mathbb{III}_1)$.

Observe that the constraint satisfaction problem Π_1 -SAT is trivial: every instance I has a solution; specifically, the map that sends every variable of I to 1. Recall that $N_{\text{NonTriv}}(I_1(\text{Cols}_3))$ is the set of all instances of Π_1 -SAT that admit only a trivial solution. Note that instances in $N_{\text{NonTriv}}(I_1(\text{Cols}_3))$ are precisely the instances of Π_1 -SAT with at most one solution. We establish the following gap property for $I_1(\text{Cols}_3)$.

Proposition 6.14. *The relation $I_1(\text{Cols}_3)$ satisfies $\text{GAP}(N_{\text{NonTriv}}, Y_{\text{SEP} \cap 2\text{-Rob}})$.*

Proof. We show there is a polynomial-time reduction taking NO instances of the problem Π_2 -SAT to instances in $N_{\text{NonTriv}}(I_1(\text{Cols}_3))$ and $Y_{\text{SEP} \cap 2\text{-Rob}}(I_2(\text{Cols}_3))$ to $Y_{\text{SEP} \cap 2\text{-Rob}}(I_1(\text{Cols}_3))$. The final result will then follow from Proposition 6.6 and Lemma 6.2.

Given an instance I of Π_2 -SAT, construct an instance $\chi(I)$ of Π_1 -SAT by replacing each constraint $\langle t = (x_1, \dots, x_8), I_2(\text{Cols}_3) \rangle$ in \mathcal{C} with $\langle t, I_1(\text{Cols}_3) \rangle$. Clearly every solution of I is a solution of $\chi(I)$. By Remark 6.7, we will also assume that there are variables \perp and \top such that every constraint tuple has \perp, \top in the final two positions.

Now assume that $I \in Y_{\text{SEP} \cap 2\text{-Rob}}(I_2(\text{Cols}_3))$. We first show that $\chi(I)$ is in $Y_{2\text{-Rob}}(I_1(\text{Cols}_3))$. Let $\{u, v\}$ be a pair of distinct elements in V and consider an assignment $\alpha: \{u, v\} \rightarrow \{0, 1\}$ that is locally compatible for $\chi(I)$.

Case 1. If α is locally compatible for I , then we can use the fact that I is in $Y_{2\text{-Rob}}(I_2(\text{Cols}_3))$ to extend α to a solution of I , which is also a solution of $\chi(I)$.

Case 2: If α is not locally compatible for I , then α satisfies $(\alpha(u), \alpha(v)) = (1, 1)$ and in this case, we can extend α to a solution of $\chi(I)$ by mapping every variable in $\chi(V)$ to 1.

To see that $\chi(I) \in Y_{\text{SEP}}(I_1(\text{Cols}_3))$, observe that solutions for I also are solutions for $\chi(I)$, and the assumption $I \in Y_{\text{SEP}}(I_2(\text{Cols}_3))$ shows that there are enough solutions to separate all pairs of variables.

We now show that if I is a NO instance of Π_2 -SAT then $\chi(I)$ has only a trivial solution. We prove the contrapositive. Suppose $\chi(I)$ has a non-trivial solution ψ in Π_1 -SAT. Hence $\psi(\perp) \neq \psi(\top)$. Then ψ is also a solution of I in Π_2 -SAT. Hence I is not a NO instance of Π_2 -SAT. \square

Corollary 6.15. *The following problems are NP-complete.*

- The 2-robust satisfiability problem $2\text{-Robust}(\text{III}_1)$
- The separation problem $\text{SEP}(\text{III}_1)$.

We establish a gap theorem for any set of relations generating Π_1 .

Theorem 6.16. *Let Γ be a constraint language on $\{0, 1\}$ such that $\langle \Gamma \rangle = \Pi_1$. Then Γ has the gap property $\text{GAP}(N_{\text{NonTriv}}, Y_{\text{SEP} \cap (2, \mathcal{G})})$, for some finite set \mathcal{G} of pp-formulae in the language of Γ , and consequently the problems $(2, \mathcal{G})\text{-Robust}(\Gamma)$ and $\text{SEP}(\Gamma)$ are NP-complete.*

Proof. The relation $I_1(\text{Cols}_3)$ is easily seen to be an irredundant weak base for the co-clone Π_1 . Thus, from Theorem 5.16, we have $I_1(\text{Cols}_3) \in \langle \Gamma \rangle_{\neq, \neq}$, and from Proposition 6.14, we have $I_1(\text{Cols}_3)$ satisfies $\text{GAP}(N_{\text{NonTriv}}, Y_{\text{SEP} \cap 2\text{-Rob}})$. The reduction given in Theorem 6.4 preserves not only the instances in $Y_{\text{SEP} \cap (2, \mathcal{F})}$ (subject to a change in \mathcal{F}), but also solutions of instances and hence it preserves instances in N_{NonTriv} . By Lemma 6.2, it then follows that Γ satisfies $\text{GAP}(N_{\text{NonTriv}}, Y_{\text{SEP} \cap (2, \mathcal{G})})$, and so $(2, \mathcal{G})\text{-Robust}(\Gamma)$ and $\text{SEP}(\Gamma)$ are NP-complete. \square

6.4. The Boolean co-clone Π_0 . Since the co-clone Π_0 is dual to Π_1 , we obtain the following results via a symmetrical argument.

In the following, let \mathbb{I}_0 be the template $\langle \{0, 1\}; I_0(\text{Cols}_3) \rangle$.

Corollary 6.17. *The following problems are NP-complete.*

- The 2-robust satisfiability problem $2\text{-Robust}(\mathbb{I}_0)$
- The separation problem $\text{SEP}(\mathbb{I}_0)$.

Theorem 6.18. *Let Γ be a constraint language on $\{0, 1\}$ such that $\langle \Gamma \rangle = \Pi_0$. Then Γ has the gap property $\text{GAP}(\text{N}_{\text{NonTriv}}, Y_{\text{SEP} \cap (2, \mathcal{G})})$, for some finite set \mathcal{G} of pp-formulae in the language of Γ , and consequently the problems $(2, \mathcal{G})\text{-Robust}(\Gamma)$ and $\text{SEP}(\Gamma)$ are NP-complete.*

6.5. The Boolean co-clone Π . By Proposition 5.12, the following relation is a weak base for the co-clone Π .

$$I(\text{Cols}_3) = \begin{bmatrix} 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}$$

We will use \mathbb{I} to denote the template $\langle \{0, 1\}; I(\text{Cols}_3) \rangle$ and $\Pi\text{-SAT}$ to denote the constraint problem $\text{CSP}(\mathbb{I})$. We remark that instances in $\text{N}_{\text{NonTriv}}(I(\text{Cols}_3))$ are exactly the instances of $\Pi\text{-SAT}$ with at most two solutions.

Proposition 6.19. *The relation $I(\text{Cols}_3)$ satisfies $\text{GAP}(\text{N}_{\text{NonTriv}}, Y_{\text{SEP} \cap 2\text{-Rob}})$.*

Proof. We show there is a polynomial-time reduction taking $\text{N}_{\text{CSP}}(I(\text{Cols}_3))$ to $\text{N}_{\text{NonTriv}}(I(\text{Cols}_3))$ and taking instances in $Y_{\text{SEP} \cap 2\text{-Rob}}(I_2(\text{Cols}_3))$ to instances in $Y_{\text{SEP} \cap 2\text{-Rob}}(I(\text{Cols}_3))$. The final result is then obtained from Proposition 6.6 and Lemma 6.2.

Given an instance I of $\Pi_2\text{-SAT}$, construct an instance $\chi_1(I)$ of $\Pi\text{-SAT}$ by replacing each constraint $\langle t = (v_1, \dots, v_8), I_2(\text{Cols}_3) \rangle$ in \mathcal{C} with $\langle t, I(\text{Cols}_3) \rangle$. Clearly every solution of I is a solution of $\chi_1(I)$. We again call on Remark 6.7 to assume that \perp and \top record the final two positions of every constraint tuple.

Assume that I is in $Y_{\text{SEP} \cap 2\text{-Rob}}(I_2(\text{Cols}_3))$. We first show that $\chi_1(I)$ is in $Y_{2\text{-Rob}}(I(\text{Cols}_3))$. Let $\{u, v\}$ be a pair of distinct elements in V and let $\alpha: \{u, v\} \rightarrow \{0, 1\}$ be an assignment that is locally compatible for $\chi_1(I)$.

Case 1: If α is locally compatible for I , then we can use the fact that $I \in Y_{2\text{-Rob}}(I_2(\text{Cols}_3))$ to extend α to a solution of I , which is also a solution of $\chi_1(I)$.

Case 2: If α is not locally compatible for I , then α satisfies $(\alpha(u), \alpha(v)) = (0, 0)$ or $(\alpha(u), \alpha(v)) = (1, 1)$. In the first case, we can extend α to a solution of $\chi_1(I)$ by mapping every variable in $\chi_1(V)$ to 0, and in the second case we extend α by mapping every variable in $\chi_1(V)$ to 1.

Now $\chi_1(I) \in Y_{\text{SEP}}(I(\text{Cols}_3))$ follows immediately from the fact that I is in $Y_{\text{SEP}}(I_2(\text{Cols}_3))$.

Since every non-trivial solution of $\chi_1(I)$ in $\Pi\text{-SAT}$ separates \perp from \top , it follows that a nontrivial solution of $\chi_1(I)$ in $\Pi\text{-SAT}$ is also a solution of I in $\Pi_2\text{-SAT}$. In the contrapositive: if I is a NO instance of $\Pi_2\text{-SAT}$, then $\chi(I)$ has only trivial solutions. \square

Corollary 6.20. *The following problems are NP-complete.*

- The 2-robust satisfiability problem 2-Robust(III).
- The separation problem SEP(III).

Theorem 6.21. *Let Γ be a constraint language on $\{0, 1\}$ such that $\langle \Gamma \rangle = \text{II}$. Then Γ has the gap property $\text{GAP}(\text{N}_{\text{NonTriv}}, \text{Y}_{\text{SEP} \cap (2, \mathcal{G})})$, for some finite set \mathcal{G} of pp-formulae in the language of Γ , and consequently the problems $(2, \mathcal{G})$ -Robust(Γ) and SEP(Γ) are **NP**-complete.*

Proof. Since $\text{I}(\text{Cols}_3)$ is an irredundant weak base for II , from Theorem 5.16, we have $\text{I}(\text{Cols}_3) \in \langle \Gamma \rangle_{\neq, \neq}$, and from Proposition 6.19, the relation $\text{I}(\text{Cols}_3)$ has gap property $\text{GAP}(\text{N}_{\text{NonTriv}}, \text{Y}_{\text{SEP} \cap 2\text{-Rob}})$. Since the reduction given in Theorem 6.4 preserves instances in $\text{Y}_{\text{SEP}} \cap \text{Y}_{(2, \mathcal{F})}$ (subject to a change in \mathcal{F}) and instances in Y_{Triv} , then from Lemma 6.2, the relation Γ satisfies $\text{GAP}(\text{N}_{\text{NonTriv}}, \text{Y}_{\text{SEP} \cap (2, \mathcal{G})})$, and so $(2, \mathcal{G})$ -Robust(Γ) and SEP(Γ) are **NP**-complete. \square

6.6. The Boolean co-clone IN. By Proposition 5.12, the following relation is a weak base for the co-clone IN.

$$\text{N}(\text{Cols}_3) = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 1 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}$$

We will use \mathbb{IN} to denote the template $\langle \{0, 1\}; \text{N}(\text{Cols}_3) \rangle$ and IN-SAT to denote the constraint problem $\text{CSP}(\mathbb{IN})$.

Proposition 6.22. *The relation $\text{N}(\text{Cols}_3)$ satisfies $\text{GAP}(\text{N}_{\text{NonTriv}}, \text{Y}_{\text{SEP} \cap 2\text{-Rob}})$.*

Proof. We show there is a polynomial-time reduction taking NO instances of the constraint problem $\text{IN}_2\text{-SAT}$ to instances in $\text{N}_{\text{NonTriv}}(\text{N}(\text{Cols}_3))$ and taking instances in $\text{Y}_{\text{SEP} \cap 2\text{-Rob}}(\text{N}_2(\text{Cols}_3))$ to instances in $\text{Y}_{\text{SEP} \cap 2\text{-Rob}}(\text{N}(\text{Cols}_3))$. The final result will then follow from Proposition 6.10 and Lemma 6.2.

Given an instance $I = (V; \{0, 1\}; \mathcal{C})$ of $\text{IN}_2\text{-SAT}$, construct an instance of $\chi_2(I)$ of IN-SAT by replacing each constraint $\langle t = (x_1, \dots, x_8), \text{N}_2(\text{Cols}_3) \rangle$ in \mathcal{C} with $\langle t, \text{N}(\text{Cols}_3) \rangle$. Clearly every solution of I is a solution of $\chi_2(I)$. We will use the assumption that \perp and \top are the only variables that appear in the final two positions of a constraint tuple, which is justified by Remark 6.11.

Now assume that I belongs to $\text{Y}_{\text{SEP} \cap 2\text{-Rob}}(\text{N}_2(\text{Cols}_3))$. We first show that $\chi_2(I) \in \text{Y}_{2\text{-Rob}}(\text{N}(\text{Cols}_3))$. Let $\{u, v\}$ be a pair of distinct elements in V and consider an assignment $\alpha: \{u, v\} \rightarrow \{0, 1\}$ that is locally compatible for $\chi_2(I)$.

Case 1. If α is locally compatible for I , then we can use the fact that I is in $\text{Y}_{2\text{-Rob}}(\text{N}_2(\text{Cols}_3))$ to extend α to a solution of I , which is also a solution of $\chi_2(I)$.

Case 2. If the assignment α is not locally compatible for I , then α satisfies $(\alpha(u), \alpha(v)) = (0, 0)$ or $(\alpha(u), \alpha(v)) = (1, 1)$. In the first case, we can extend α to a solution of $\chi_2(I)$ by sending every variable in V to 0 and in the second, we extend by sending every variable in V to 1.

Now $\chi_2(I) \in \text{Y}_{\text{SEP}}(\text{N}(\text{Cols}_3))$ follows from the fact that $I \in \text{Y}_{\text{SEP}}(\text{N}_2(\text{Cols}_3))$: we use the separating solutions of I , which are also solutions of $\chi_2(I)$, to separate distinct pairs of elements of $\chi_2(I)$.

Now assume that $\chi_2(I)$ has a non-trivial solution ψ in IN-SAT. Then ψ separates \perp from \top , so is also a solution of I in IN_2 -SAT. In the contrapositive: if I is a NO instance of IN_2 -SAT, then $\chi_2(I)$ has only trivial solutions for IN-SAT. \square

Corollary 6.23. *The following problems are NP-complete.*

- The 2-robust satisfiability problem $2\text{-Robust}(\mathbb{IN})$.
- The separation problem $\text{SEP}(\mathbb{IN})$.

Theorem 6.24. *Let Γ be a constraint language on $\{0, 1\}$ such that $\langle \Gamma \rangle = \text{IN}$. Then Γ has the gap property $\text{GAP}(\text{N}_{\text{NonTriv}}, \text{Y}_{\text{SEP} \cap (2, \mathcal{G})})$, for some finite set \mathcal{G} of ppositive formulæ in the language of Γ , and consequently the problems $(2, \mathcal{G})\text{-Robust}(\Gamma)$ and $\text{SEP}(\Gamma)$ are NP-complete.*

Proof. Since $\text{N}(\text{Cols}_3)$ is an irredundant weak base for the co-clone IN, from Theorem 5.16, we have $\text{N}(\text{Cols}_3) \in \langle \Gamma \rangle_{\exists, \neq}$, and from Proposition 6.22, the relation $\text{N}(\text{Cols}_3)$ has gap property $\text{GAP}(\text{N}_{\text{NonTriv}}, \text{Y}_{\text{SEP} \cap 2\text{-Rob}})$. Since the reduction given in Theorem 6.4 preserves instances in $\text{Y}_{\text{SEP} \cap (2, \mathcal{F})}$ (subject to a change in \mathcal{F}) and instances in $\text{N}_{\text{NonTriv}}$, then by Lemma 6.2, that Γ satisfies $\text{GAP}(\text{N}_{\text{NonTriv}}, \text{Y}_{\text{SEP} \cap (2, \mathcal{G})})$, for some finite set \mathcal{G} of pp-formulæ, and consequently $(2, \mathcal{G})\text{-Robust}(\Gamma)$ and $\text{SEP}(\Gamma)$ are NP-complete. \square

7. TOWARDS A TRICHOTOMY: PROVING TRACTABILITY

In this section, we investigate tractability for $\text{SEP}(\Gamma)$ and $(2, \mathcal{F})\text{-Robust}(\Gamma)$ over constraint languages Γ on $\{0, 1\}$. We establish a theorem that covers all cases that are solvable in polynomial-time, see Theorem 7.2 below.

In the following it useful to recall that a finite relational structure (template) is a core if all of its unary polymorphisms are bijective.

The following theorem appears in [21, Proposition 3.2]; strictly speaking, the result in [21] holds for k -robust satisfiability, however the proof is applicable to the broader setting of (k, \mathcal{F}) -robust satisfiability. We list only the properties required to establish our main results.

Theorem 7.1. [21, Proposition 3.2] *Let Γ be a constraint language on a finite set A and let \mathbb{A} be the corresponding template. Consider $\mathbb{A}_0 = \langle A; \Gamma \cup \{(a) \mid a \in A\} \rangle$. There is a logspace Turing reduction from the following computational problems to $\text{CSP}(\mathbb{A}_0)$.*

- $(k, \mathcal{F})\text{-Robust}(\mathbb{A})$, for any $k \in \mathbb{N}$ and any finite set \mathcal{F} of pp-formulae in the language of Γ .
- $\text{SEP}(\mathbb{A})$.

Recall that to a core structure \mathbb{A} we can add all singleton unary relations without increasing the complexity of $\text{CSP}(\mathbb{A})$, see [8, Theorem 4.7]. Thus, it follows from Theorem 7.1, that in the case where \mathbb{A} is a core, each of the problems listed is reducible to $\text{CSP}(\mathbb{A})$.

Theorem 7.2. *Let Γ be a constraint language on $\{0, 1\}$. If $\text{IN} \not\subseteq \langle \Gamma \rangle$, then the computational problems $\text{SEP}(\Gamma)$ and $(2, \mathcal{F})\text{-Robust}(\Gamma)$ are solvable in polynomial-time.*

Proof. When $\text{IN} \not\subseteq \langle \Gamma \rangle$, it follows from Post's co-clone lattice (see Figure 1), that $\text{IN} \not\subseteq \langle \Gamma \cup \{(0), (1)\} \rangle$ and then it is known that the constraint problem

$\text{CSP}(\Gamma \cup \{(0), (1)\})$ is tractable; this can be found in Schaefer's original argument for example; see [33, Lemma 4.1]. Then from Theorem 7.1, the problems $\text{SEP}(\Gamma)$ and $(2, \mathcal{F})\text{-Robust}(\Gamma)$ are solvable in polynomial-time. \square

8. AN ALGEBRAIC APPROACH

The algebraic analysis of polymorphisms has been a very powerful tool in the study of CSP complexity; for the main development of this approach see [25, 26, 27, 28]. In particular, the existence of polymorphisms satisfying certain equations has been shown to ensure tractability (see [24] for example) and can be used to identify appropriate algorithmic methods for solvability [3, 16]. Conversely, the absence of polymorphisms satisfying certain equations ensures intractability [29]. A first step in developing the algebraic theory was recognising that CSP complexity is preserved under the operator $\langle - \rangle$; this goes back to Schaefer [33]. A second step was connecting the algebra of polymorphisms to fundamental concepts in universal algebra: first variety theory, then the deeper analysis of tame congruence theory; much of the universal algebraic analysis is built on the work of [8]. For a comprehensive overview of the algebraic approach see [2, 23].

We have seen that for some computational problems the $\langle - \rangle$ operator appears too coarse. When the weaker operator $\langle - \rangle_{\mathcal{F}}$ can be substituted, conventional polymorphism analysis can be replaced by the analysis of *partial* polymorphisms. In this section, we examine which aspects of the second step in the algebraic analysis of CSP complexity can be recovered for problems amenable to the $\langle - \rangle_{\mathcal{F}}$ operator. In particular, we investigate how the computational complexity of SEP , $(k, \mathcal{F})\text{-Robust}$ and other variants of the CSP relate to variety-like properties of the partial polymorphism algebra and establish theorems that are analogous to some of the main contributions in Bulatov et al. [8, Theorems 4.7, 5.2 & 5.4].

We first prove a more general version of Theorem 6.4. The addition of equality presents complications for $(2, \mathcal{F})\text{-robust}$ satisfiability, and we require the use of a gap property to carry through the reduction. The proof for SEP is similar, however a further slight variation of the proof is necessary. We will say that an equality constraint $\langle (x_i, x_j), =_A \rangle$ is non-trivial if $i \neq j$, otherwise it is trivial.

Theorem 8.1. *Let A be a non-empty finite set and let Γ be a constraint language over A . Let R be any finite subset of $\langle \Gamma \rangle_{\mathcal{F}}$ and let \mathcal{F} be a finite set of pp-formulae in the language of R . Then there is a polynomial-time Karp reduction from $\text{SEP}(R)$ to $\text{SEP}(\Gamma)$. Moreover, if R has gap property $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP} \cap (k, \mathcal{F})})$, then Γ satisfies $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP} \cap (k, \mathcal{G})})$, for some finite set \mathcal{G} of pp-formulae in the language of Γ .*

Proof. First fix some NO instance J of $\text{SEP}(\Gamma)$, ensuring that if $\text{CSP}(\Gamma)$ is non-trivial, then J is a NO instance of $\text{CSP}(\Gamma)$ (which is automatically a NO instance of $\text{SEP}(\Gamma)$).

We now make a small adjustment to the proof of Lemma 6.3. Given an R -instance I , we construct a Γ -instance I' . Again, each constraint $\langle (v_1, \dots, v_\ell), r^A \rangle \in \mathcal{C}$ is to be replaced by constraints built from the conjuncts in the formula defining r from Γ . However, now there may be equality conjuncts.

Case 1. If there is a non-trivial equality constraint $\langle (x_i, x_j), =_A \rangle$ introduced, then I was a NO instance of $\text{SEP}(R)$ as no satisfying solution can separate x_i and x_j . Thus, for I' we may output J .

Case 2. If all equality constraints constructed are trivial, then we may omit these from the conjuncts. All of the conjuncts are now equality-free and so we are in the

situation encountered in Lemma 6.3: the solutions of the instance I are identical to those of I' . Hence I is a YES instance of $\text{SEP}(R)$ if and only if I is a YES instance of $\text{SEP}(\Gamma)$. As this is trivially true in Case 1 also, we have the desired reduction.

Assume that R satisfies $\text{GAP}(\text{NCSP}, Y_{\text{SEP} \cap (k, \mathcal{F})})$. It follows that $\text{CSP}(R)$ is nontrivial and so $\text{CSP}(\Gamma)$ is non-trivial as $\text{CSP}(R)$ reduces to $\text{CSP}(\Gamma)$. Assume I is a NO instance of $\text{CSP}(R)$. If Case 1 applies, then $I' = J$ is NO instance of $\text{CSP}(\Gamma)$, as required. If Case 2 applies, then I' cannot have a solution. So I' is again a NO instance of $\text{CSP}(\Gamma)$.

Now assume $I \in Y_{\text{SEP} \cap (k, \mathcal{F})}$. Then Case 2 applies and the solutions of the instance I are identical to those of I' . We are then essentially in the equality-free case considered in (2) and (3) of Theorem 6.4. Using the set \mathcal{G} and argument given there, it follows that I' is in $Y_{\text{SEP} \cap (k, \mathcal{G})}$. The final result is obtained by Lemma 6.2. \square

Results of Bulatov, Jeavons and Krokhin provide the fundamental connection between CSP complexity and the tame-congruence theoretic analysis of polymorphism algebras; see [8, Theorems 5.2, 5.4] or [29, Theorem 2.1] for a refined version. We now prove a variation of this result, which we establish through a series of polynomial-time reductions, presented in the form of lemmas. The constructions used for substructures (Lemma 8.5) and homomorphisms (Lemma 8.4) are based on those in [8, 29], given in the standard CSP setting. In the case of taking substructures, SEP and $(2, \mathcal{F})$ -Robust carry through using the standard construction (subject to a change in the local compatibility condition \mathcal{F}). The homomorphism case however requires proper amendment, including the addition of extra relations and non-trivial usage of the gap property.

For partial algebras there are a number of reasonable ways to define the concepts of subalgebra and homomorphism; see Grätzer [20] for example. We take the definition of subalgebra coinciding with the model-theoretic notion of substructure and take the weakest (that is, most general) definition of homomorphism of the standard definitions.

Definition 8.2. Let F be a set of finitary partial operation symbols. The partial algebra $\mathbf{B} = \langle B; F^{\mathbf{B}} \rangle$ is a *subalgebra* of the partial algebra $\mathbf{A} = \langle A; F^{\mathbf{A}} \rangle$, if $B \subseteq A$ and for each n -ary $f \in F$, we have $\text{dom}(f^{\mathbf{B}}) = \text{dom}(f^{\mathbf{A}}) \cap B^n$, and $f^{\mathbf{B}}$ agrees with $f^{\mathbf{A}}$ on this set.

Definition 8.3. Let F be a set of finitary partial operation symbols and let $\mathbf{A} = \langle A; F^{\mathbf{A}} \rangle$ and $\mathbf{B} = \langle B; F^{\mathbf{B}} \rangle$ be partial algebras. A map $\varphi: A \rightarrow B$ is a *homomorphism* if, for every n -ary $f \in F$ and $(a_1, \dots, a_n) \in \text{dom}(f^{\mathbf{A}})$, we have $(\varphi(a_1), \dots, \varphi(a_n)) \in \text{dom}(f^{\mathbf{B}})$ and $\varphi(f^{\mathbf{A}}(a_1, \dots, a_n)) = f^{\mathbf{B}}(\varphi(a_1), \dots, \varphi(a_n))$.

We say that \mathbf{B} is a *homomorphic image* of \mathbf{A} if there exists a surjective homomorphism $\varphi: \mathbf{A} \twoheadrightarrow \mathbf{B}$. Note that homomorphic images are not uniquely determined by the map $\varphi: A \rightarrow B$.

In the following preliminary lemmas, it may be useful to recall that, for any non-empty set A and set F of partial operations on A , a k -ary relation r on A is in $\text{Inv}(F)$ precisely when r is a subuniverse of the partial algebra $\langle A; F \rangle^k$.

Lemma 8.4. Let F be a set of finitary partial operation symbols, let $\mathbf{A} = \langle A; F^{\mathbf{A}} \rangle$ and $\mathbf{B} = \langle B; F^{\mathbf{B}} \rangle$ be finite partial algebras and let $\varphi: \mathbf{A} \twoheadrightarrow \mathbf{B}$ be a surjective

homomorphism. Let Γ be a finite subset of $\text{Inv}(F^{\mathbf{B}})$, let \mathcal{F} be a finite set of pp-formulae in the language of Γ and let $k \in \mathbb{N}$. Then for some finite set $\Gamma' \subseteq \text{Inv}(F^{\mathbf{A}})$ there is a polynomial-time computable function reducing

- (1) $\text{CSP}(\langle B; \Gamma \rangle)$ to $\text{CSP}(\langle A; \Gamma' \rangle)$,
- (2) $(k, \mathcal{F})\text{-Robust}(\langle B; \Gamma \rangle)$ to $(k, \mathcal{G})\text{-Robust}(\langle A; \Gamma' \rangle)$, for some finite set \mathcal{G} of pp-formulae in the language of Γ' ,

and taking $Y_{\text{SEP}}(\Gamma)$ to $Y_{\text{SEP}}(\Gamma')$.

Proof. We use a slight variation to the construction given in the proof of [8, Theorem 5.4]. We begin by finding a subset Γ' of $\text{Inv}(F^{\mathbf{A}})$: for each k -ary relation r in $\text{Inv}(F^{\mathbf{B}})$, construct the following set

$$\varphi^{\leftarrow}(r) := \{(a_1, \dots, a_k) \mid (\varphi(a_1), \dots, \varphi(a_k)) \in r\} \subseteq A^k,$$

and let $\Gamma' = \{\varphi^{\leftarrow}(r) \mid r \in \Gamma\} \cup \{\ker(\varphi)\}$. Certainly Γ' is finite and it is an easy exercise to check that $\ker(\varphi)$ is a subuniverse of $\langle A; F^{\mathbf{A}} \rangle^2$ and $\varphi^{\leftarrow}(r) \in \text{Inv}(F^{\mathbf{A}})$, for each r in $\text{Inv}(F^{\mathbf{B}})$.

Given an instance $I = (V; B; \mathcal{C})$ of $\text{CSP}(\langle B; \Gamma \rangle)$, we construct an instance $I' = (V; A; \mathcal{C}')$ of $\text{CSP}(\langle A; \Gamma' \rangle)$ by replacing each constraint $\langle s; r \rangle \in \mathcal{C}$ with $\langle s; \varphi^{\leftarrow}(r) \rangle$. Thus, $\mathcal{C}' = \{\langle s; \varphi^{\leftarrow}(r) \rangle \mid \langle s; r \rangle \in \mathcal{C}\}$.

The proof for (1) is identical to the proof of [8, Theorem 5.4]. For (2), we begin by finding a set \mathcal{G} of pp-formulae in the language of Γ' . For each $\rho(w_1, \dots, w_n) \in \mathcal{F}$ we construct a pp-formula $\rho_{\Gamma'}(w_1, \dots, w_n)$ in the language of Γ' by replacing every occurrence of $r \in \Gamma$ with $\varphi^{\leftarrow}(r) \in \Gamma'$ and every occurrence of $=_B$ with $\ker(\varphi)$. Observe that $\rho_{\Gamma'}(a_1, \dots, a_n)$ is true in $\langle A; \Gamma' \rangle$ if and only if $\rho(\varphi(a_1), \dots, \varphi(a_n))$ is true in $\langle B; \Gamma \rangle$, for any $(a_1, \dots, a_n) \in A^n$ and $\rho(w_1, \dots, w_n) \in \mathcal{F}$. Now assume that $I = (V; B; \mathcal{C})$ is a YES instance of $(k, \mathcal{F})\text{-Robust}(\langle B; \Gamma \rangle)$ and consider an assignment $\alpha: S \rightarrow A$ with $|S| = k$. If α is \mathcal{G} compatible for I' , then $\varphi \circ \alpha: S \rightarrow B$ is \mathcal{F} -compatible for I . We can extend $\varphi \circ \alpha$ to a solution ϕ of I . Then any function $f: V \rightarrow A$ such that $\varphi \circ f(v) = \phi(v)$, for all $v \in V$, is a solution of I' that extends α . For the converse direction, begin by assuming that $I' = (V; A; \mathcal{C}')$ is a YES instance of $(k, \mathcal{G})\text{-Robust}(\langle A; \Gamma' \rangle)$ and consider an assignment $\beta: S \rightarrow B$ with $|S| = k$ that is \mathcal{F} -compatible for I . Then any $g: S \rightarrow A$ such that $\varphi \circ g(v) = \beta(v)$, for all $v \in V$ is \mathcal{G} -compatible for I' . We can extend g to a solution ϕ of I' . Then $\varphi \circ \phi: V \rightarrow B$ is a solution of I that extends β .

To prove (3), assume that I is a YES instance of $\text{SEP}(\langle B; \Gamma \rangle)$. Let $\{v_1, v_2\}$ be a pair of distinct variables in V and let $\phi: V \rightarrow B$ be a separating solution of I for the pair $\{v_1, v_2\}$ guaranteed by the assumption. Then any function $\alpha: V \rightarrow A$ such that $\varphi \circ \alpha(v) = \phi(v)$, for all $v \in V$, is a separating solution of I' for the pair $\{v_1, v_2\}$. Hence I' is a YES instance of $\text{SEP}(\langle A; \Gamma' \rangle)$. \square

Lemma 8.5. *Let F be a set of finitary partial operation symbols, let $\mathbf{A} = \langle A; F^{\mathbf{A}} \rangle$ be a finite partial algebra and let $\mathbf{B} = \langle B; F^{\mathbf{B}} \rangle$ be a subalgebra of \mathbf{A} . Let Γ be a finite subset of $\text{Inv}(F^{\mathbf{B}})$, let \mathcal{F} a finite subset of pp-formulae in the language of Γ and let $k \in \mathbb{N}$. Then for some finite set Γ' of $\text{Inv}(F^{\mathbf{A}})$ there is a polynomial-time reduction from*

- (1) $\text{CSP}(\langle B; \Gamma \rangle)$ to $\text{CSP}(\langle A; \Gamma' \rangle)$,
- (2) $\text{SEP}(\langle B; \Gamma \rangle)$ to $\text{SEP}(\langle A; \Gamma' \rangle)$,
- (3) $(k, \mathcal{F})\text{-Robust}(\langle B; \Gamma \rangle)$ to $(k, \mathcal{G})\text{-Robust}(\langle A; \Gamma' \rangle)$, for some finite set \mathcal{G} of pp-formulae in the language of Γ' .

Proof. We first show that $\text{Inv}(F^{\mathbf{B}}) \subseteq \text{Inv}(F^{\mathbf{A}})$. Let r be a k -ary relation on the set B and assume that r is invariant under $F^{\mathbf{B}}$. We will show that r is invariant under $F^{\mathbf{A}}$. Let $f^{\mathbf{A}} \in F^{\mathbf{A}}$ be n -ary, and let $a_1 = (a_{11}, \dots, a_{1n}), a_2 = (a_{21}, \dots, a_{2n}), \dots, a_k = (a_{k1}, \dots, a_{kn})$ be tuples in $\text{dom}(f^{\mathbf{A}}) \subseteq A^n$. Now assume that for all $i \in \{1, \dots, n\}$, we have $(a_{1i}, a_{2i}, \dots, a_{ki}) \in r$. Since r is a relation on B , we have $\{a_1, \dots, a_k\} \subseteq \text{dom}(f^{\mathbf{B}}) = \text{dom}(f^{\mathbf{A}}) \cap B^n$ and since $f^{\mathbf{A}} = f^{\mathbf{B}}$ on this set, it follows that

$$(f^{\mathbf{A}}(a_1), f^{\mathbf{A}}(a_2), \dots, f^{\mathbf{A}}(a_k)) = (f^{\mathbf{B}}(a_1), f^{\mathbf{B}}(a_2), \dots, f^{\mathbf{B}}(a_k)) \in r,$$

and hence $r \in \text{Inv}(F^{\mathbf{A}})$. Now since \mathbf{B} is a subalgebra of \mathbf{A} , the set B , viewed as a unary relation on A , belongs to $\text{Inv}(F^{\mathbf{A}})$. So let $\Gamma' = \Gamma \cup \{B\} \subseteq \text{Inv}(F^{\mathbf{A}})$.

The following construction is borrowed from [8, Theorem 5.2] and gives a reduction from $\text{CSP}(\langle B; \Gamma \rangle)$ to $\text{CSP}(\langle A; \Gamma' \rangle)$, proving (1). Given an instance $I = (V; B; \mathcal{C})$ of $\text{CSP}(\langle B; \Gamma \rangle)$, we construct an instance of $I' = (V; A; \mathcal{C}')$ of $\text{CSP}(\langle A; \Gamma' \rangle)$ in the following way. For each $v \in V$, add the constraint $\langle (v), B \rangle$ to \mathcal{C} . So $\mathcal{C}' = \mathcal{C} \cup \{\langle (v), B \rangle \mid v \in V\}$. Note that the solutions of $I = (V; B; \mathcal{C})$ are identical to those of $I' = (V; A; \mathcal{C}')$. It follows that the separating solutions of I are precisely the separating solutions of I' . Hence I is a YES instance of $\text{SEP}(\langle B; \Gamma \rangle)$ if and only if I' is a YES instance of $\text{SEP}(\langle A; \Gamma' \rangle)$. This proves (2).

For (3), assume that I is a YES instance of $(k, \mathcal{F})\text{-Robust}(\langle B; \Gamma \rangle)$ and let $\mathcal{G} = \mathcal{F} \cup \{(x) \in B\}$. Every \mathcal{G} -compatible assignment on k variables of I' preserves \mathcal{F} -constraints into B as $\langle (v), B \rangle$ is a constraint of I' , for all $v \in V$. Thus, every \mathcal{G} -compatible assignment on k variables for I' is an \mathcal{F} -compatible assignment for I , and conversely. The reduction from $(k, \mathcal{F})\text{-Robust}(\Gamma)$ to $(k, \mathcal{G})\text{-Robust}(\Gamma')$ then follows from the fact that the solutions of $I = (V; B; \mathcal{C})$ are exactly the solutions of $I' = (V; A; \mathcal{C}')$. \square

We are now ready to prove the first main result of this section. For any partial algebra \mathbf{A} we let $\text{HS}(\mathbf{A})$ be the smallest class of partial algebras in the same signature closed under the formation of homomorphic images (H) and subalgebras (S).

Theorem 8.6 (HS Gap Theorem). *Let $\mathbb{A} = \langle A; R \rangle$ and $\mathbb{B} = \langle B; S \rangle$ be templates, let \mathcal{F} be a finite set of pp-formulae in the language of R and let $k \in \mathbb{N}$. If S has gap property $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP} \cap (k, \mathcal{F})})$ and there exist partial algebras $\mathbf{A} = \langle A; F^{\mathbf{A}} \rangle$ and $\mathbf{B} = \langle B; F^{\mathbf{B}} \rangle$ such that*

- (1) $F^{\mathbf{B}} \subseteq \text{pPol}(S)$,
- (2) $\mathbf{B} \in \text{HS}(\mathbf{A})$, and
- (3) $\text{pPol}(R) \subseteq F^{\mathbf{A}}$,

then R has gap property $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP} \cap (k, \mathcal{G})})$.

Proof. The result is established by carrying the gap property through items (1), (2) and (3). This requires showing there is a series of polynomial-time reductions preserving NO instances of CSP and instances that are simultaneously YES instances of SEP and $(k, \mathcal{F})\text{-Robust}$.

From (1), we have $\text{Inv}(\text{pPol}(S)) \subseteq \text{Inv}(F^{\mathbf{B}})$ and so $S \subseteq \langle S \rangle_{\overline{\mathcal{F}}} \subseteq \text{Inv}(F^{\mathbf{B}})$. From (2), there is a surjective homomorphism $\varphi: \mathbf{C} \twoheadrightarrow \mathbf{B}$, where $\mathbf{C} = \langle C; F^{\mathbf{C}} \rangle$ is a subalgebra of \mathbf{A} . Thus, by Lemma 8.4, there is a finite subset T of $\text{Inv}(F^{\mathbf{C}})$ and a polynomial-time reduction taking instances in $\text{N}_{\text{CSP}}(S)$ to $\text{N}_{\text{CSP}}(T)$, and instances in $\text{Y}_{\text{SEP} \cap (k, \mathcal{F})}(S)$ to $\text{Y}_{\text{SEP} \cap (k, \mathcal{F}')(T)}$, for some finite set \mathcal{F}' of pp-formulae

in the language of T . Hence T satisfies $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP} \cap (k, \mathcal{F}')})$. Now since \mathbf{C} is a subalgebra of \mathbf{A} , by Lemma 8.5, there is a finite set U of $\text{Inv}(F^{\mathbf{A}})$ and a polynomial-time reduction taking instances in $\text{N}_{\text{CSP}}(T)$ to $\text{N}_{\text{CSP}}(U)$, and instances in $\text{Y}_{\text{SEP} \cap (k, \mathcal{F}')} (T)$ to $\text{Y}_{\text{SEP} \cap (k, \mathcal{G}')} (U)$, for some finite set \mathcal{G}' of pp-formulae in the language of U . Hence U satisfies $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP} \cap (k, \mathcal{G}')})$. From (3), it follows that $\text{Inv}(F^{\mathbf{A}}) \subseteq \text{Inv}(\text{pPol}(R^{\mathbf{A}})) = \langle R \rangle_{\overline{\mathcal{F}}}$. Thus, U is a finite subset of $\langle R \rangle_{\overline{\mathcal{F}}}$. By Theorem 8.1, it then follows that R has gap property $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP} \cap (k, \mathcal{G})})$, for some finite set \mathcal{G} of pp-formulae in the language of R . \square

Remark 8.7. When $\mathbf{A} \neq \mathbf{B}$, Theorem 8.6 lifts the complete classification given by the Gap Trichotomy Theorem on Boolean domains to many problems on templates with non-Boolean domains. With further effort, direct products (P) can be incorporated into condition (2) of Theorem 8.6, however the argument is beyond the scope of the present article, and will appear in subsequent work. For now we establish a weaker version.

We first introduce the following equivalence relations.

Definition 8.8. Let A be a non-empty set and $\ell \in \mathbb{N}$. For any pair of elements $a_1 = (a_{1,0}, \dots, a_{1,\ell-1}), a_2 = (a_{2,0}, \dots, a_{2,\ell-1})$ in A^ℓ and for every $p, q \in \{0, \dots, \ell-1\}$, we define

$$a_1 \equiv_{pq} a_2 \text{ if } a_{1,p} = a_{2,q},$$

For any $\ell \in \mathbb{N}$, there are ℓ^2 equivalence relations of the form \equiv_{pq} , for $p, q \in \{0, \dots, \ell-1\}$.

The basis for the following construction was inspired by [29, Lemma 2.4], however the reduction requires significant modification with the main addition being a pre-processing “local reflection” procedure.

Lemma 8.9. *Let F be a set of finitary partial operation symbols, let $\mathbf{A} = \langle A; F^{\mathbf{A}} \rangle$ be a finite partial algebra and consider $\mathbf{A}^\ell = \langle A; F^{\mathbf{A}^\ell} \rangle$. Let Γ be a finite subset of $\text{Inv}(F^{\mathbf{A}^\ell})$ and let \mathcal{F} be a finite set of pp-formulae in the language of Γ . Then for some finite set $\Gamma' \subseteq \text{Inv}(F^{\mathbf{A}})$ there is a polynomial-time function taking*

- (1) $\text{N}_{\text{CSP}}(\Gamma)$ to $\text{N}_{\text{CSP}}(\Gamma')$,
- (2) and $\text{Y}_{(2, \mathcal{F})}(\Gamma)$ to $\text{Y}_{\text{SEP}}(\Gamma')$.

Proof. We begin by fixing the subset Γ' of $\text{Inv}(F^{\mathbf{A}})$. A k -ary relation r on A^ℓ can be considered as a ℓk -ary relation r^b on A , by replacing each k -tuple of ℓ -tuples

$$((a_{1,0}, \dots, a_{1,\ell-1}), \dots, (a_{k,0}, \dots, a_{k,\ell-1}))$$

by the ℓk -ary tuple

$$(a_{1,0}, \dots, a_{1,\ell-1}, \dots, a_{k,0}, \dots, a_{k,\ell-1}).$$

In particular, if $\mathbf{A} = \langle A; F^{\mathbf{A}} \rangle$ is a partial algebra and r is a subuniverse of the direct power $(\mathbf{A}^\ell)^k$, then r^b is the corresponding subuniverse of $\mathbf{A}^{\ell k}$. Hence $r^b \in \text{Inv}(F^{\mathbf{A}})$ and so let $\Gamma' = \{r^b \mid r \in \Gamma\}$.

The reduction now proceeds in two steps. Given an instance $I = (V; A^\ell; \mathcal{C})$ of $\text{CSP}(\Gamma)$ where $V = \{v_i \mid i \leq |V|\}$, we begin by constructing an instance $I^b = (V^b; A; \mathcal{C}^b)$ of $\text{CSP}(\Gamma' \cup \{=_A\})$ in the following way.

- (i) Let $V^b = \{v_{i,p} \mid v_i \in V \text{ and } p \in \{0, \dots, \ell-1\}\}$, and

(ii) replace each constraint $\langle (v_1, \dots, v_k), r \rangle$ in \mathcal{C} with

$$\langle (v_{1,0}, \dots, v_{1,\ell-1}, \dots, v_{k,0}, \dots, v_{k,\ell-1}), r^b \rangle.$$

At this point in the construction, we denote the instance by I_0 . To complete the construction of I^b , we apply the following “local reflection” procedure to I . For every $i \leq |V|$, select some number $i' \neq i$ with $i' \leq |V|$. For every $v_i \in V$ and $p, q \in \{0, \dots, \ell - 1\}$ with $p \neq q$,

- if every \mathcal{F} -compatible assignment $\alpha: \{v_i, v_{i'}\} \rightarrow A^\ell$ maps (v_i, v_i) into \equiv_{pq} , we add the constraint $\langle (v_{i,p}, v_{i,q}), =_A \rangle$,

and for every distinct pair $\{v_i, v_j\}$ in V and $p, q \in \{0, \dots, \ell - 1\}$,

- if every \mathcal{F} -compatible assignment $\alpha: \{v_i, v_j\} \rightarrow A^\ell$ maps (v_i, v_j) into \equiv_{pq} , we add $\langle (v_{i,p}, v_{j,q}), =_A \rangle$.

We now construct an instance $I^b = (V^b; A; \mathcal{C}^b)$ of $\text{CSP}(\Gamma')$, by eliminating all constraints involving $=_A$ in the usual way: fix some arbitrary total order $<$ on V^b and let \equiv_A denote the reflexive transitive closure of the relation $=_A$ on V^b . For any $v_{i,p}$, let $\sigma(v_{i,p})$ denote the $<$ -smallest element of the \equiv_A -class of $v_{i,p}$. Then I^b is simply I^b , but with each appearance of $v_{i,p} \in V^b$ replaced by $\sigma(v_{i,p})$ (and all $=_A$ constraints removed).

For (1), we prove the contrapositive. If φ is a solution of I^b in $\text{CSP}(\Gamma')$, then $\varphi \circ \sigma: V^b \rightarrow A$ is a solution of I^b and therefore is also a solution of I_0 . The map $\psi: V \rightarrow A^\ell$ defined by $\psi(v_i) = (\varphi(\sigma(v_{i,0})), \dots, \varphi(\sigma(v_{i,\ell-1})))$, for all $v_i \in V$, is a solution of I in $\text{CSP}(\Gamma)$.

For (2), assume that I is a YES instance of $(2, \mathcal{F})\text{-Robust}(\Gamma)$. We will show that I^b is a YES instance of $\text{SEP}(\Gamma')$. Let $\{v_{i,p}, v_{j,q}\}$ be a distinct pair of variables in V^b . We consider two cases.

Case 1. If $i = j$ and $p \neq q$. Then there was some \mathcal{F} -compatible assignment $\alpha: \{v_i, v_{i'}\} \rightarrow A^\ell$ mapping (v_i, v_i) into $A^\ell \setminus \equiv_{pq}$. Hence $\pi_p(\alpha(v_i)) \neq \pi_q(\alpha(v_i))$. We can extend α to a full solution φ of I in $\text{CSP}(\Gamma)$ as we are assuming that I is $(2, \mathcal{F})$ -robustly satisfiable. Then the map $\varphi^b: V^b \rightarrow A$ defined by $\varphi^b(v_{st}) = \pi_t(\varphi(v_s))$, for all $s \leq |V|$ and $t \in \{0, \dots, \ell - 1\}$, is a solution of I^b that separates the pair $\{v_{i,p}, v_{i,q}\}$.

Case 2. If $i \neq j$ and $p, q \in \{0, \dots, \ell - 1\}$. Since $v_{i,p}$ and $v_{j,q}$ were not identified by the above construction, there must have been some \mathcal{F} -compatible assignment $\alpha: \{v_i, v_j\} \rightarrow A^\ell$ mapping (v_i, v_j) into $A^\ell \setminus \equiv_{pq}$. Hence $\pi_p(\alpha(v_i)) \neq \pi_q(\alpha(v_j))$. We can extend α to a solution φ of I in $\text{CSP}(\Gamma)$. Then the map $\varphi^b: V^b \rightarrow A$ defined by $\varphi^b(v_{st}) = \pi_t(\varphi(v_s))$, for all $s \leq |V|$ and $t \in \{0, \dots, \ell\}$, is a solution of I^b that separates the pair $\{v_{i,p}, v_{j,q}\}$. It follows from Case 1 and 2 that I^b is a YES instance of $\text{SEP}(\Gamma')$. \square

Theorem 8.10 (HSP Gap Theorem). *Let $\mathbb{A} = \langle A; R \rangle$ and $\mathbb{B} = \langle B; S \rangle$ be templates and let \mathcal{F} be a finite set of pp-formulae in the language of R . If S has the gap property $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP} \cap (2, \mathcal{F})})$ and there exist partial algebras $\mathbf{A} = \langle A; F^{\mathbf{A}} \rangle$ and $\mathbf{B} = \langle B; F^{\mathbf{B}} \rangle$ such that*

- (1) $F^{\mathbf{B}} \subseteq \text{pPol}(S)$,
- (2) $\mathbf{B} \in \text{HSP}(\mathbf{A})$, and
- (3) $\text{pPol}(R) \subseteq F^{\mathbf{A}}$,

then R satisfies $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP}})$.

Proof. The argument follows the proof of Theorem 8.6 up to and including the application of Lemma 8.5, except now \mathbf{C} is a subalgebra of \mathbf{A}^ℓ , for some $\ell \in \mathbb{N}$. We then apply Lemma 8.9 before continuing on to the application of condition (3). The conclusion that R has gap property $\text{GAP}(\text{N}_{\text{CSP}}, \text{Y}_{\text{SEP}})$ is then established from the proof of Theorem 8.1 (in contrast to Theorem 8.1 directly). \square

We now observe that HS theorems can be obtained for other variants of the constraint satisfaction problem. The following *equivalence problem* has been studied in Böhler, Hemaspaandra, Reith and Vollmer [7] and [34], and the *implication problem* has been studied in [34].

Equivalence problem $\text{EQUIV}(\mathbb{A})$ over template $\mathbb{A} = \langle A; \Gamma^{\mathbb{A}} \rangle$.

Instance: a pair of $\Gamma^{\mathbb{A}}$ -instances $(I_1 = (V; A; \mathcal{C}_1), I_2 = (V; A; \mathcal{C}_2))$.

Question: for every assignment φ , is it true that φ is a solution of I_1 if and only if φ is a solution of I_2 ?

Implication problem $\text{IMPL}(\mathbb{A})$ over template $\mathbb{A} = \langle A; \Gamma^{\mathbb{A}} \rangle$.

Instance: a pair of $\Gamma^{\mathbb{A}}$ -instances $(I_1 = (V; A; \mathcal{C}_1), I_2 = (V; A; \mathcal{C}_2))$.

Question: for every assignment φ , is it true that if φ is a solution of I_1 then φ is a solution of I_2 ?

As observed in [34, §6.2], the two problems are closely related, in fact they are equivalent under polynomial-time Turing reductions: a pair (I_1, I_2) is a YES instance of $\text{IMPL}(\mathbb{A})$ if and only if $(I_1, (V; A; \mathcal{C}_1 \cup \mathcal{C}_2))$ is a YES instance of $\text{EQUIV}(\mathbb{A})$, while (I_1, I_2) is a YES instance of $\text{EQUIV}(\mathbb{A})$ if and only if both (I_1, I_2) and (I_2, I_1) are YES instances of $\text{IMPL}(\mathbb{A})$.

In the case of Boolean domains, both EQUIV and IMPL have been shown to experience \mathbf{P} versus CoNP dichotomies with identical boundaries to those established for SEP and $(2, \mathcal{F})$ -Robust; see [7] for the original proof and [34] for a simplified proof using clone-theoretic techniques. The HS theorem for EQUIV and IMPL (Theorem 8.13 below) lifts the complete classifications given by these dichotomies on Boolean domains to templates with non-Boolean domains. This is particularly useful for EQUIV and IMPL as it is unknown if the operator $\langle - \rangle$ can be applied in the case of non-Boolean domains. Even the operator $\langle - \rangle_{\overline{\neq}}$ is unknown to be compatible with EQUIV in the non-Boolean case. The weaker operator $\langle - \rangle_{\overline{\neq}, \neq}$ is known to be applicable (see Lemma 8.12 below), but this imposes the additional assumption of irredundancy.

The proof of the HS theorem for EQUIV and IMPL relies on the following two results. It may be useful to recall the notion of irredundancy given in Definition 5.14.

Lemma 8.11. [35, Proposition 3.11] *Let R be a constraint language over a non-empty finite set A and let S be a finite subset of $\langle R \rangle_{\overline{\neq}}$. If S is irredundant, then $S \subseteq \langle R \rangle_{\overline{\neq}, \neq}$.*

Lemma 8.12. [34, Proposition 2.2] *Let R be a constraint language over a non-empty finite set A and let $S \subseteq \langle R \rangle_{\overline{\neq}, \neq}$. Then there is polynomial-time reduction from $\text{EQUIV}(S)$ to $\text{EQUIV}(R)$ and from $\text{IMPL}(S)$ to $\text{IMPL}(R)$.*

Theorem 8.13 (HS Theorem for EQUIV and IMPL). *Let $\mathbb{A} = \langle A; R \rangle$ and $\mathbb{B} = \langle B; S \rangle$ be templates. If S is irredundant and there exist partial algebras $\mathbf{A} = \langle A; F^{\mathbf{A}} \rangle$ and $\mathbf{B} = \langle B; F^{\mathbf{B}} \rangle$ such that*

- (1) $F^{\mathbf{B}} \subseteq \text{pPol}(S)$,
- (2) $\mathbf{B} \in \text{HS}(\mathbf{A})$, and
- (3) $\text{pPol}(R) \subseteq F^{\mathbf{A}}$,

then there is a polynomial-time reduction from $\text{EQUIV}(S)$ to $\text{EQUIV}(R)$ and from $\text{IMPL}(S)$ to $\text{IMPL}(R)$.

Proof. We prove the result only for IMPL as the proof for EQUIV is similar.

Condition (1) implies that $\text{Inv}(\text{pPol}(S)) \subseteq \text{Inv}(F^{\mathbf{B}})$ and so $S \subseteq \langle S \rangle_{\overline{\mathcal{A}}} \subseteq \text{Inv}(F^{\mathbf{B}})$. Condition (2) gives a surjective homomorphism $\varphi: \mathbf{C} \rightarrow \mathbf{B}$, where $\mathbf{C} = \langle C; F^{\mathbf{C}} \rangle$ is a subalgebra of \mathbf{A} .

We first show that for some $T \subseteq \text{Inv}(F^{\mathbf{C}})$, the problem $\text{IMPL}(\langle B; S \rangle)$ is reducible to $\text{IMPL}(\langle C; T \rangle)$ in polynomial-time. We use the construction given in the proof of Lemma 8.4, except now there is no need for the inclusion of the relation $\ker(\varphi)$. Let $T = \{\varphi^{\leftarrow}(r) \mid r \in S\} \subseteq \text{Inv}(F^{\mathbf{C}})$. Given an instance

$$I = (I_1 = (V; B; \mathcal{C}_1), I_2 = (V; B; \mathcal{C}_1))$$

of $\text{IMPL}(\langle B; S \rangle)$, we construct an instance

$$I' = (I'_1 = (V; C; \mathcal{C}'_1), I'_2 = (V; C; \mathcal{C}'_2))$$

of $\text{IMPL}(\langle C; T \rangle)$ by letting $\mathcal{C}'_i = \{\langle s; \varphi^{\leftarrow}(r) \rangle \mid \langle s; r \rangle \in \mathcal{C}_i\}$, for each $i \in \{1, 2\}$. Now assume that I is a YES instance of $\text{IMPL}(\langle B; S \rangle)$. Consider a solution $\psi: V \rightarrow C$ of I'_1 . Then $\varphi \circ \psi: V \rightarrow B$ is a solution of I_1 and is therefore a solution of I_2 . Then any function $\mu: V \rightarrow C$ satisfying $\varphi \circ \mu = \varphi \circ \psi$ is a solution of I'_2 . In particular, ψ is a solution of I'_2 . Conversely, assume that I' is a YES instance of $\text{IMPL}(\langle C; T \rangle)$ and consider a solution $\phi: V \rightarrow B$ of I_1 . Let $\nu: V \rightarrow C$ be any function satisfying $\varphi \circ \nu = \phi$. Then ν is a solution of I'_1 and therefore also a solution of I'_2 . It then follows that $\varphi \circ \nu = \phi$ is a solution of I_2 , as required.

We now show that for some $U \subseteq \text{Inv}(F^{\mathbf{A}})$, there is a polynomial-time reduction from $\text{IMPL}(\langle C; T \rangle)$ to $\text{IMPL}(\langle A; U \rangle)$. In the degenerate case that $\mathbf{C} = \mathbf{A}$, we let $U = T$. Otherwise we use the construction given in the proof of Lemma 8.5. Let $U = T \cup \{C\}$. Given an instance $I = (I_1, I_2)$ of $\text{IMPL}(\langle C; T \rangle)$, we construct an instance $I' = (I'_1, I'_2)$ of $\text{IMPL}(\langle A; U \rangle)$ in the following way: for each $i \in \{1, 2\}$, we let $\mathcal{C}'_i = \mathcal{C} \cup \{(v) \in B \mid v \in V\}$. The reduction immediately follows from the fact that the solutions of I_i are identical to those of I'_i , for each $i \in \{1, 2\}$.

Now from (3), we have that $\text{Inv}(F^{\mathbf{A}}) \subseteq \text{Inv}(\text{pPol}(R^{\mathbf{A}})) = \langle R \rangle_{\overline{\mathcal{A}}}$. Thus, U is a finite subset of $\langle R \rangle_{\overline{\mathcal{A}}}$. Observe that the constructions applied to the constraint languages S and T above preserve irredundancy, and so U is also irredundant. By Lemma 8.11, it follows that U is a finite subset of $\langle R \rangle_{\overline{\mathcal{A}}, \neq}$. Then Lemma 8.12 gives the final reduction from $\text{IMPL}(\langle A; U \rangle)$ to $\text{IMPL}(\langle A; R \rangle)$. \square

A further useful simplification in the standard CSP setting has been the restriction to so-called idempotent polymorphisms; see [8, Theorem 4.7]. A partial polymorphism $f: \text{dom}(f) \rightarrow \mathbf{A}$ is *idempotent* if $f(a, \dots, a) = a$ for every $a \in A$ for which $(a, \dots, a) \in \text{dom}(f)$. As a final result we show that when analysing the complexity of SEP problems we may restrict to idempotent partial polymorphisms.

We first fix some useful notation. For a template $\mathbf{A} = \langle A; \Gamma \rangle$, let V_A be a copy of A given by $\{v_a \mid a \in A\}$ (to be treated as variables), and let $\text{diag}(\mathbf{A})$ denote the equality-free positive atomic diagram of \mathbf{A} :

$$\text{diag}(\mathbf{A}) = \{ \langle (v_{a_1}, \dots, v_{a_n}), r \rangle \mid (a_1, \dots, a_n) \in r, \text{ where } r \text{ is an } n\text{-ary relation in } \Gamma \}.$$

Let Γ_{Con} denote the union of Γ with the set of all singleton unary relations $\{(a)\}$ for $a \in A$. The proof of the following theorem is based on examples described by Jackson [22].

Theorem 8.14. *Let Γ be a constraint language over a finite non-empty set A . If $\mathbb{A} = \langle A; \Gamma \rangle$ is a core relational structure, then $\text{SEP}(\Gamma)$ is equivalent to $\text{SEP}(\Gamma_{\text{Con}})$, with respect to a polynomial-time Turing reduction.*

Proof. Every instance of $\text{SEP}(\Gamma)$ is an instance of $\text{SEP}(\Gamma_{\text{Con}})$ as $\Gamma \subseteq \Gamma_{\text{Con}}$, and is equivalently a YES instance (or NO instance) of both problems. Hence $\text{SEP}(\Gamma)$ reduces to $\text{SEP}(\Gamma_{\text{Con}})$ in constant-time.

We now show there is a polynomial-time Turing reduction from $\text{SEP}(\Gamma_{\text{Con}})$ to $\text{SEP}(\Gamma)$. For a given instance $I = (V; A; \mathcal{C})$ of $\text{SEP}(\Gamma_{\text{Con}})$, we begin by constructing a new instance $I_A = (V'; A; \mathcal{C}')$ of $\text{SEP}(\Gamma_{\text{Con}})$ in the following way: first let V' be the disjoint union $V \cup V_A$ and construct \mathcal{C}' by replacing every constraint $\langle v, \{(a)\} \rangle$ in \mathcal{C} with $\langle (v, v_a), =_A \rangle$, and then taking the disjoint union with $\text{diag}(\mathbb{A})$.

We will show there is a polynomial-time computable family \mathcal{F} of Γ -instances satisfying the property that I is a YES instance of $\text{SEP}(\Gamma_{\text{Con}})$ if and only if there exists $I' \in \mathcal{F}$ such that I' is a YES instance of $\text{SEP}(\Gamma)$.

Case 1. If there exists $a \in A$ and constraints $\{\langle (v_1, v_a), =_A \rangle, \langle (v_2, v_a), =_A \rangle\} \subseteq \mathcal{C}'$ with $v_1 \neq v_2$, then I is automatically a NO instance of $\text{SEP}(\Gamma_{\text{Con}})$ as any satisfying solution of I must identify v_1 and v_2 . In this case, we let \mathcal{F} be the set containing a single fixed NO instance of $\text{SEP}(\Gamma)$.

Case 2. For every $a \in A$ at most one variable in V is constrained to $\{(a)\}$. In this case there may be some $a \in A$ for which $\{(a)\}$ does not appear in any constraint in \mathcal{C} . Let U be the set of all such “unused” elements in A and let V_U be the set of all variables in V that are not constrained to any singleton unary relation. We first examine the number of ways to extend \mathcal{C}' to include singleton constraints from U without falling back into Case 1. Each member of \mathcal{F} will be constructed from such an extension. For every subset $S \subseteq U$ and injective function ι from U into V_U , let $I^\iota = (V'; A; \mathcal{C}' \cup \{\langle (\iota(a), v_a), =_A \rangle \mid a \in S\})$. Let $m := |V_U|$ and let n_i be the number of i -element subsets of U , for each $i \in \{0, \dots, |U|\}$. Let $k = \max\{n_i \mid i = 0, \dots, |U|\} \leq 2^{|A|}$. The number of injective functions from an i element set to V_U is the i -permutation ${}^m P_i \leq m^{|A|}$. Then the total number of instances included in \mathcal{F} is

$$\sum_{i \in \{0, \dots, |U|\}} {}^m P_i \times n_i \leq (m+1) \times m^{|A|} \times 2^{|A|}.$$

As $|A|$ is bounded by a constant and $m \leq |V|$, there is a uniform polynomial bound on the number of instances included in \mathcal{F} (in fact, the total number of instances in \mathcal{F} is in $O(|V|^{|A|+1})$).

Consider an extension $I^\iota = (V', A, \mathcal{C}' \cup \{\langle (\iota(a), v_a), =_A \rangle \mid a \in S\})$ of I . Apply the usual procedure to I^ι for eliminating all constraints involving the equality relation; the resulting variables will be denoted V'' . When $\iota(a) = v$, it is notationally convenient to use both v and v_a to denote the same point of V'' after this identification: this is justified by the observation that in this construction there are no sequences of equalities of length more than one, so that distinct elements of V , as well as distinct elements of V_A , remain distinct after the identification. In particular, this abuse of notation allows us to consider both V_A and V as (overlapping) subsets of V' . Let \mathcal{F} be the set of these constructed instances.

Assume there is some $I' \in \mathcal{F}$ such that I' is a YES instance of $\text{SEP}(\Gamma)$. So for every pair $\{v_1, v_2\}$ of distinct variables in V'' there is a separating solution φ for that pair. Since φ preserves $\text{diag}(\mathbb{A})$, the restriction $\varphi|_{V_A}$ determines an automorphism α of \mathbb{A} , by way of $\alpha: a \mapsto b$ if $\varphi(v_a) = b$. Then the map $\alpha^{-1} \circ \varphi|_V$ is a solution of I separating the pair $\{v_1, v_2\}$. Hence I is a YES instance of $\text{SEP}(\Gamma_{\text{Con}})$.

Now assume that I is a YES instance of $\text{SEP}(\Gamma_{\text{Con}})$. We will show there exists some $I' \in \mathcal{F}$ such that I' is a YES instance of $\text{SEP}(\Gamma)$. We first apply the following reflection procedure to I_A . For each $v \in V_U$, if every homomorphism from I to $\langle A; \Gamma_{\text{Con}} \rangle$ maps $v \in V_U$ to some $u \in U$, then we add the constraint $\langle (v, v_u), =_A \rangle$ to \mathcal{C}' . Note that for distinct $v, v' \in V$ we cannot have both $\langle (v, v_u), =_A \rangle$ and $\langle (v', v_u), =_A \rangle$ added, as we are assuming there is a solution separating v from v' . Thus the instance constructed in this way coincides with one of the extensions from Case 2. Find the $I' \in \mathcal{F}$ corresponding to this extension. For every pair of distinct variables $\{v_1, v_2\}$ in V there is a solution ψ of I separating v_1 and v_2 . Extend the separating solution ψ for $\{v_1, v_2\}$ to a separating solution of I' by sending v_a to a . Hence I' is a YES instance of $\text{SEP}(\Gamma)$. \square

Definition 8.15. Let F be a set of finitary partial operation symbols and let $\mathbf{A} = \langle A; F^{\mathbf{A}} \rangle$ be a finite partial algebra. We define the partial algebra

$$\mathbf{A}_0 := \langle A; \text{id}([F^{\mathbf{A}}]_p) \rangle,$$

where $\text{id}([F^{\mathbf{A}}]_p)$ consists of all idempotent operations from the strong partial clone generated by $F^{\mathbf{A}}$. We refer to \mathbf{A}_0 as the *full idempotent reduct* of \mathbf{A} .

Notice that $\text{Inv}(\text{id}([F^{\mathbf{A}}]_p)) = \text{Inv}(F^{\mathbf{A}})_{\text{Con}}$. The next result is an immediate consequence of Theorem 8.14

Corollary 8.16. Let $\mathbf{A} = \langle A; F^{\mathbf{A}} \rangle$ be a finite partial algebra and consider the full idempotent reduct \mathbf{A}_0 of \mathbf{A} . If $\mathbb{A} = \langle A; \text{Inv}(F^{\mathbf{A}}) \rangle$ is a core relational structure, then $\text{SEP}(\text{Inv}(F^{\mathbf{A}}))$ is equivalent to $\text{SEP}(\text{Inv}(F^{\mathbf{A}})_{\text{Con}})$, with respect to a polynomial-time Turing reduction.

9. CONCLUSION

The results obtained demonstrate that gap properties can prove very powerful in the study of the complexity of constraint problems. In particular, the Gap Trichotomy Theorem provides dichotomies for an entire family of constraint problems on Boolean domains. Gap properties also facilitated an algebraic approach to the study of constraint problems amenable to partial polymorphism analysis: validity of most of the reductions relied intrinsically on the assumption of a gap property. A possible future direction for work might be to extend this analysis to a tame-congruence theoretic approach. This would require a reworking of tame-congruence theory to adjust for partial algebras. Could such an analysis lead to algebraic dichotomy conjectures for new variants of the constraint satisfaction problem?

In future work, our focus will extend beyond **P** versus **NP** dichotomies to a finer-grained analysis of the complexity of the reductions given above. While we did not develop these directions in the present article, arguments in the style of Larose and Tesson [29] can be used to perform all of the reductions in log-space, and most using first order reductions. Details for tightening the reductions will be given in a subsequent paper, where we examine natural extensions of the SEP and $(2, \mathcal{F})$ -Robust conditions and their degree of compatibility with the operator $\langle - \rangle$.

A gap dichotomy of the style given in the present article is believed to be true for these problems, however more technical machinery is required.

Acknowledgements. The author is indebted to Marcel Jackson for many useful suggestions and discussions, as well as her PhD supervisors Brian Davey and Tomasz Kowalski, for their guidance and feedback.

REFERENCES

- [1] Samson Abramsky, Georg Gottlob, and Phokion G. Kolaitis. Robust constraint satisfaction and local hidden variables in quantum mechanics. In Francesca Rossi, editor, *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*. IJCAI/AAAI, 2013.
- [2] Libor Barto. The constraint satisfaction problem and universal algebra. *Bull. Symb. Log.*, 21(3):319–337, 2015.
- [3] Libor Barto and Marcin Kozik. Constraint satisfaction problems solvable by local consistency methods. *J. ACM*, 61(1):3:1–3:19, 2014.
- [4] Libor Barto, Marcin Kozik, and Todd Niven. The CSP dichotomy holds for digraphs with no sources and no sinks (A positive answer to a conjecture of bang-jensen and hell). *SIAM J. Comput.*, 38(5):1782–1802, 2009.
- [5] Adam Beacham. The complexity of problems without backbones. Master’s thesis, University of Alberta, Canada, 2000.
- [6] Adam Beacham and Joseph C. Culberson. On the complexity of unfrozen problems. *Discrete Applied Mathematics*, 153(1-3):3–24, 2005.
- [7] Elmar Böhler, Edith Hemaspaandra, Steffen Reith, and Heribert Vollmer. Equivalence and isomorphism for boolean constraint satisfaction. In Julian C. Bradfield, editor, *Computer Science Logic, 16th International Workshop, CSL 2002, 11th Annual Conference of the EACSL, Edinburgh, Scotland, UK, September 22-25, 2002, Proceedings*, volume 2471 of *Lecture Notes in Computer Science*, pages 412–426. Springer, 2002.
- [8] Andrei Bulatov, Peter Jeavons, and Andrei Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM J. Comput.*, 34(3):720–742, 2005.
- [9] Andrei A. Bulatov. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *J. ACM*, 53(1):66–120, January 2006.
- [10] Andrei A. Bulatov. Complexity of conservative constraint satisfaction problems. *ACM Trans. Comput. Logic*, 12(4):24:1–24:66, July 2011.
- [11] Andrei A. Bulatov. The complexity of the counting constraint satisfaction problem. *J. ACM*, 60(5):34, 2013.
- [12] Nadia Creignou and Jean-Jacques Hébrard. On generating all solutions of generalized satisfiability problems. *ITA*, 31(6):499–511, 1997.
- [13] Nadia Creignou, Phokion G. Kolaitis, and Bruno Zanuttini. Structure identification of boolean relations and plain bases for co-clones. *J. Comput. Syst. Sci.*, 74(7):1103–1115, 2008.
- [14] Joseph C. Culberson and Ian P. Gent. Frozen development in graph coloring. *Theor. Comput. Sci.*, 265(1-2):227–264, 2001.
- [15] Brian A. Davey, Jane G. Pitkethly, and Ross Willard. The lattice of alter egos. *Internat. J. Algebra Comput.*, 22(1):1250007, 36, 2012.
- [16] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM J. Comput.*, 28(1):57–104, 1998.
- [17] David Geiger. Closed systems of functions and predicates. *Pacific J. Math.*, 27:95–100, 1968.
- [18] Oded Goldreich. On promise problems: A survey. In *Theoretical Computer Science, Essays in Memory of Shimon Even*, volume 3895 of *Lecture Notes in Computer Science*, pages 254–290. Springer, 2006.
- [19] Georg Gottlob. On minimal constraint networks. In Jimmy Ho-Man Lee, editor, *Principles and Practice of Constraint Programming - CP 2011 - 17th International Conference, CP 2011, Perugia, Italy, September 12-16, 2011. Proceedings*, volume 6876 of *Lecture Notes in Computer Science*, pages 325–339. Springer, 2011.

- [20] George Grätzer. *Universal algebra*. Springer, New York, second edition, 2008. With appendices by Grätzer, Bjarni Jónsson, Walter Taylor, Robert W. Quackenbush, Günter H. Wenzel, and Grätzer and W. A. Lampe.
- [21] M. Jackson. Flexible constraint satisfiability and a problem in semigroup theory. arXiv:1512.03127 [math.LO], 2015.
- [22] Marcel Jackson. Personal conversation, May 2016.
- [23] Marcel Jackson, Tomasz Kowalski, and Todd Niven. Complexity and polymorphisms for digraph constraint problems under some basic constructions. To appear in *Internat. J. Algebra Comput.* arXiv:1304.4986 [math.CO], 2013.
- [24] Peter Jeavons. On the algebraic structure of combinatorial problems. *Theor. Comput. Sci.*, 200(1-2):185–204, 1998.
- [25] Peter Jeavons, David A. Cohen, and Martin C. Cooper. Constraints, consistency and closure. *Artif. Intell.*, 101(1-2):251–265, 1998.
- [26] Peter Jeavons, David A. Cohen, and Marc Gyssens. A unifying framework for tractable constraints. In *Principles and Practice of Constraint Programming - CP'95, First International Conference, CP'95, Cassis, France, September 19-22, 1995, Proceedings*, pages 276–291, 1995.
- [27] Peter Jeavons, David A. Cohen, and Marc Gyssens. Closure properties of constraints. *J. ACM*, 44(4):527–548, 1997.
- [28] Peter Jeavons, David A. Cohen, and Justin Pearson. Constraints and universal algebra. *Ann. Math. Artif. Intell.*, 24(1-4):51–67, 1998.
- [29] Benoit Larose and Pascal Tesson. Universal algebra and hardness results for constraint satisfaction problems. *Theor. Comput. Sci.*, 410(18):1629–1647, 2009.
- [30] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. Determining computational complexity from characteristic phase transitions. *Nature*, 400:133–137, 1998.
- [31] E. L. Post. *The two-valued iterative systems of mathematical logic*, volume no. 5 of *Annals of Mathematics studies*. Princeton University Press, Princeton, 1941.
- [32] B. A. Romov. The algebras of partial functions and their invariants. *Cybernetics*, 17:157–167, 1981.
- [33] Thomas J. Schaefer. The complexity of satisfiability problems. In Richard J. Lipton, Walter A. Burkhard, Walter J. Savitch, Emily P. Friedman, and Alfred V. Aho, editors, *Proceedings of the 10th Annual ACM Symposium on Theory of Computing, May 1-3, 1978, San Diego, California, USA*, pages 216–226. ACM, 1978.
- [34] Henning Schnoor and Ilka Schnoor. Partial polymorphisms and constraint satisfaction problems. In Nadia Creignou, Phokion G. Kolaitis, and Heribert Vollmer, editors, *Complexity of Constraints - An Overview of Current Research Themes [Result of a Dagstuhl Seminar]*., volume 5250 of *Lecture Notes in Computer Science*, pages 229–254. Springer, 2008.
- [35] Ilka Schnoor. *The weak base method for constraint satisfaction*. PhD thesis, University of Hanover, 2008.
- [36] Karsten Schölzel. Dichotomy on intervals of strong partial Boolean clones. *Algebra Universalis*, 73(3-4):347–368, 2015.
- [37] Edward P. K. Tsang. *Foundations of constraint satisfaction*. Computation in cognitive science. Academic Press, 1993.

DEPARTMENT OF MATHEMATICS AND STATISTICS, LA TROBE UNIVERSITY, VICTORIA 3086, AUSTRALIA

E-mail address: leham@students.latrobe.edu.au